arXiv:1612.09212v1 [cs.AI] 29 Dec 2016

TECHNISCHE UNIVERSITÄT
KAISERSLAUTERN

Fraunhofer
IDMT

Technische Universität
Kaiserlautern
Fachbereich Informatik

Bachelor Thesis

zur Erlangung des akademischen Grades
Bachelor of Science

Thema:	A hybrid approach to supervised machine learning for algorithmic melody composition
Autor:	Rouven Bauer <r_bauer11@cs.uni-kl.de>
Version vom:	23. Oktober 2016
1. Betreuer:	Prof. Dr. Prof. h.c. Andreas Dengel
2. Betreuer:	Dr. rer. nat. Stephan Baumann
externe Betreuung:	Dr.-Ing. Estefanía Cano Cerón

Acknowledgements

This work would not have been possible without the great help and supervision of Dr.-Ing. Estefanía Cano Cerón at Fraunhofer IDMT in Ilemnau. Thanks for all the time, inspiring discussions and ideas. Thanks also to the supervisors at DFKI Kaiserslautern Dr. Stephan Baumann and Prof. Anreas Dengel as well as all other people giving their support in manners like administrative support, guidance, proof-reading and motivation.

Abstract

English:
In this work we present an algorithm for composing monophonic melodies similar in style to those of a given, phrase annotated, sample of melodies. For implementation, a hybrid approach incorporating parametric Markov models of higher order and a contour concept of phrases is used. This work is based on the master thesis of Thayabaran Kathiresan (2015). An online listening test conducted shows that enhancing a pure Markov model with musically relevant context, like count and planed melody contour, improves the result significantly.

German:
In dieser Arbeit wird ein Algorithmus präsentiert, der einstimmige Melodien komponiert. Diese sind stilistisch ähnlich zu Beispielmelodien, deren Phrasengrenzen zuvor annotiert wurden. Dazu wird ein Hybridansatz verfolgt, der parametrische Markow-Modelle höherer Ordnung, sowie ein Konturenkonzept implementiert. Diese Arbeit basiert auf der Masterarbeit von Thayabaran Kathiresan (2015). Ein Onlineumfrage hat ergeben, dass das erweitern eines reinen Markow-Modells mit musikalisch relevantem Kontext, wie Zählzeit und geplanter Melodiekontur, die Ergebnisse signifikant verbessert.

Contents

Definitions 7

1. Introduction 9
1.1. Motivation and Goal . . . 9

2. Previous Work 10
2.1. State of the Art . . . 10
2.1.1. Generating Melodies with Grammars . . . 10
2.1.2. Generating Melodies with other Methods of symbolic AI . . . 11
2.1.3. Generating Melodies with Evolutionary Algorithms . . . 12
2.1.4. Generating Melodies with Neural Networks . . . 12
2.1.5. Generating Melodies with Markov Chains . . . 13
2.2. Related Work . . . 14
2.2.1. "Automatic Melody Generation" by Thayabaran Kathiresan . . . 14

3. Musical Terms 15
3.1. Key and Mode . . . 15
3.2. Beats, Bars, and Time Signature . . . 16
3.3. Count and Off-beat . . . 16

4. Proposed Method 17
4.1. Learning . . . 19
4.1.1. MTC-Dataset . . . 19
4.1.2. Parametric Markov models . . . 20
4.1.3. Contours . . . 22
4.2. Phrase Composition . . . 24
4.2.1. Joining Pitch and Rhythm Generation . . . 24
4.2.2. Ending Constraints . . . 27
4.2.3. Following Contours . . . 28
4.2.4. Backtracking . . . 28

5. Results 29
5.1. Generating the Baseline . . . 29
5.2. Conducted Listening Test . . . 29
5.3. Results of the Listening Test . . . 30
5.4. Interpretation of the Results . . . 33

6. Conclusion and Future Work 35

Bibliography 37

Appendix 40

List of Figures

1. Taxonomy of methods of previous work 10
2. Pleasantness of Kathiresan's results 15
3. Intervals of major and minor scale 16
4. Note durations 16
5. Counting $\frac{4}{4}$ bars 17
6. Workflow of proposed method 18
7. Workflow of contour learning 25
8. Clustering results of contour learning 26
9. Survey result: age 31
10. Survey result: self-estimated musicality 31
11. Survey result: gender, instrument playing, origin and field of activity . 32
12. Survey result: liking of melody types 33
13. Correlation of phrase liking and knowing it 34

List of Tables

1. Transition matrix example 14
2. Confusion matrix of Kathiresan's results 14
3. Mode distribution of MTC-FS 20
4. Time signature distribution of MTC-FS 21
5. Melody types presented to the subjects 30
6. Raw melody responses 32
7. Percental melody responses 34

Definitions

In this thesis we use the following mathematical notation:

- Sets
  - $|X|$ for a set X is it's cardinality (i.e. the number of elements contained).
  - $[a,\, b] \subset \mathbb{R},\ a, b \in \mathbb{R}$ is an interval of the real numbers. It contains all elements $x \in \mathbb{R} \colon a \leq x \leq b$.
  - $(a,\, b) \subset \mathbb{R},\ a, b \in \mathbb{R}$ is an interval of the real numbers. It contains all elements $x \in \mathbb{R} \colon a < x < b$.
  - Combinations like $[a,\, b)$ and $(a,\, b]$ are possible as well.
  - $\{a,\, \ldots,\, b\} \subset \mathbb{Z},\ a, b \in \mathbb{Z}$ is an interval of integers. It contains all elements $x \in \mathbb{Z} \colon a \leq x \leq b$.
- Functions
  - $f \colon X \mapsto Y$ is a function (or mapping), called f, from the set X to the set Y.
  - $\varphi_{\mu,\sigma^2}(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$ is the Gaussian distribution with μ as mean and σ^2 as standard deviation.
  - $\mathrm{FFT}(f)$ is the fast Fourier transform of the function f.
  - $\mathrm{FFT}^{-1}(f)$ is the inverse fast Fourier transform of the function f.
- Vectors
  - $\vec{x}$ is a vector.
  - For an n-dimensional vector $\vec{x} \in \mathbb{R}^n$ we define $\dim(\vec{x}) = n$ as the dimension of the vector.
  - $\|\vec{x}\| = \sqrt{\sum_i \vec{x}_i^2}$ is the Euclidean length of a vector $\vec{x}$.
  - $\langle \vec{x} \rangle = \begin{cases} \frac{\vec{x}}{\|\vec{x}\|} & \|\vec{x}\| \neq 0 \\ \vec{0} & \text{else} \end{cases}$ is a vector normalized by it's Euclidean length.
  - $\vec{a} \bigodot \vec{b} = \vec{c}$ with $\vec{c}_i = \vec{a}_i \cdot \vec{b}_i$ is the element wise multiplication of two vectors.
  - $\vec{b} = \mathrm{accsum}(\vec{a})$, the accumulative sum defined element-wise as $\vec{b_i} = \sum_{j \leq i} \vec{a_j}$.

1. Introduction

Artificial intelligence has made huge progress over the last decades. One topic still fascinating and demanding for further research is *computational creativity* (sometimes also referred to as *artificial creativity*). Either attempting to understand what human creativity is and how it works by developing computational models for it or driven by the attempt to outperform humans at one of the tasks they are probably still better in than computers. Composing music is one of those creative tasks. Music is a concept known to all cultural groups around the world. It plays an important role in many social contexts like religion and plain entertainment.

The idea of automatic composition is an old one. W. A. Mozart proposed such a method in 1798 [Moz98]. Since there were no computers at his time he utilized a pair of dice to randomly recombine a set a bars to compose *Countrydances*. There were as many as $1.85 \cdot 10^{17}$ different pieces that could be composed with his method. It is an example of automatically composing a melody with harmonization arranged for piano all at once. It is possible though, to subdivide the composing process into several task. Examples are *rhythm generation*, *harmonization* and *melody composition*. This work focuses on the latter one.

1.1. Motivation and Goal

One definition of music is "organized sound". This definition is a very broad one. We want to focus on western style music, more exactly melodies, that can be written down in a traditional staff. It's worth noting that written music still leaves a lot of interpretive freedom, like slight variations in pitch or time as well as agogics, for each musician that may play it. This music interpretation process is a different topic that goes beyond the scope of our work. An example of such work was proposed by Flossmann et al. [FGW13]. For a broad overview of the topic the interested reader is referred to the survey by Kirke and Miranda [KM09].

What our system tries to achieve is composing monophonic melodies similar to those of a given training set of melodies. It utilizes higher order Markov models and other statistical models to do so. We decided to follow this approach as this work builds upon the master thesis of Kathiresan [Kat15] which incorporates Markov models as well. The source code can be found on GitHub (`https://github.com/roba91/melody-composer`).

A system like this could find applications in generating melodies for practicing an instrument or music theoretical tasks like harmonization with as many melodies as a student pleases. Since it's more fun and therefore more motivating to play or analyze melodies that are pleasant and don't feel unnatural or artificial, we tried to improve the

original system by Kathiresan (see Section 2.2.1) in a way to compose more human-like melodies than it was able to create before.

2. Previous Work

2.1. State of the Art

There are several ways to structure the previous work on the topic of algorithmic composition. Here, we will stick to the taxonomy as proposed by Fernández and Vico [FV13]. Their survey paper provides a great overview of the state of the art and is recommended to the reader for a more in depth insight.

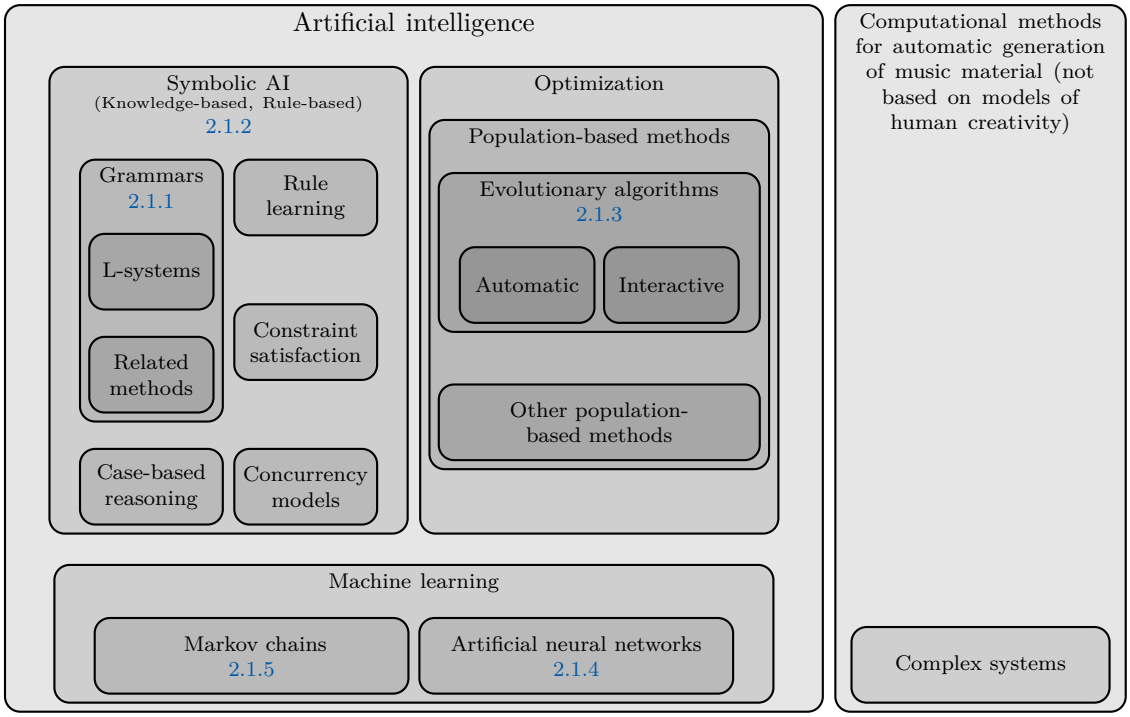

Figure 1: Taxonomy of approaches in previous work of computer composed music as proposed by Fernández and Vico [FV13]. Not all fields are covered in our work because we want to focus on the major approaches and those relevant to ours.

It is worth mentioning that it is very difficult in general to clearly categorize all previous works by their approach. This is because many of them are hybrid approaches (like the one proposed by us), utilizing more the one *artificial intelligence* (AI) approach. Nevertheless, we categorized them by their most dominating approach.

2.1.1. Generating Melodies with Grammars

Lindenmayer Systems (L-systems) and *Chomsky grammars* have in common that they are an intuitive choice for algorithmic composition. Music is very often repetitive and self-similar. This self-similarity can be found on multiple scales in music. In groups of

just a few notes up to bars, phrases, whole compositions (as in movements) or sometimes even across pieces. This self-similarity corresponds to the hierarchical structure of grammars.

Grammars are defined by a set of *production rules* that are applied recursively in so called *derivations* to substitute parts of a symbol chain with other symbols. L-systems work similarly, but in every derivation step, all possible rules are applied simultaneously. Such a generated symbol chain must then be mapped to a musical object. Often, this is a simple translation from the grammar alphabet to aspects of a composition, such as pitches, durations, or chords. Moorer [Moo72] and Langston [Lan89] proposed systems that apply these methods.

One important step in grammar driven composition is to generate the rules that define the grammar. There are two basic strategies for this: 1) One can either derive the rules manually from music theory as Rader [Rad74] did, or 2) the recently more often chosen strategy, is to abstract from the musical structure of a training dataset to obtain those rules as Gillick et al. [GTK09] did.

2.1.2. Generating Melodies with other Methods of symbolic AI

Rule-based systems are another intuitive choice for algorithmic composition. Music theory offers a lot of rules on how to compose music, arrange it and other tasks. Some examples of such rules that are well-founded and well-defined in music theory are counter point [MM95] and basso continuo [WK07].

The rule set used for composing is commonly a static one during execution; however, there are some proposals that use a combination of static and dynamically learned rules [Sch93]. The Literature also provides systems that derive their rules completely from training data [Spa99], or combine rule templates with training data [MM95].

Another approach from the field of symbolic AI is the formulation of a set of constraints. The musical task is then accomplished by solving the resulting *constraint satisfaction problem* (CSP). A lot of work using CSPs was done on solving classical problems like harmonization or counterpoint composition [RP98]. But other aspects of musical composition have been addressed as well. Two consecutive CSPs can be used to compose music following a storyboard modeling the mood of the composition as a function of time as shown by Zimmermann [Zim01]. The first CSP generates a harmony progression from the storyboard while the second one then composes a four-part harmonization from the first's output. Another application of CSPs was proposed by Laruson and Kuuskankare [LK00]. They used constraints to ensure human playability of a melody in respect to the fingering for example.

Case-based reasoning (CBR) is a further approach to implement rules of composition tasks [PGMC97]. CBR is an approach where the system starts off with a set of cases for which the solution to the problem is known. If a solution for a new (unknown) case

is needed, the system tries to find a similar case and to modify the solution for it in order to solve the new case. This case-solution pair can then be added to the set of known cases.

2.1.3. Generating Melodies with Evolutionary Algorithms

Evolutionary algorithms in general start with an initial set (called *population*) of candidate solutions (called *individuals*, *creatures* or *phenotypes*) to a problem. In iterative steps, new populations (called *generations*) are generated so that they are better candidate solutions. Therefore, a *fitness function* must be defined that describes the quality of the candidates. The fitter individuals of the latest generation are selected and modified by so called *cross-overs* and *mutations* which are combinations and random changes of the individuals to form the next generation.

This approach has been used in two major forms: Espí et al. [ELPS+07] and Jensen [Jen11] are recent examples of works using a well-defined fitness function. Nevertheless, it is very hard, if not impossible, to find a good mathematical model of aesthetic perception. Thus, approaches with human interaction replacing or complementing a defined fitness function by manually rating individuals were proposed. MacCallum [MMBL12] published one of many examples.

2.1.4. Generating Melodies with Neural Networks

Artificial neural networks, called *neural networks* for simplicity in this text, are a machine learning approach inspired by nature. These networks are composed of *layers* built out of so called *neurons*. In feed-forward networks, neurons take some or all (depending on the type of layer) neuron outputs of the previous layer as inputs. The inputs are then multiplied by *weights* (usually different for all inputs), summed up and mapped with an *activation function* (most simple activation function is the signum function). The weights are what is altered when a neural network is trained [Kri07, DHS12]. *Recurrent neural networks* like *LSTMs*[1] have more complicated link structures between *neurons* that allow the network for having an internal state similar to a memory [HS97]. This comes in handy when one does not want to compose a melody as a whole but sequentially.

Neural networks as a tool in AI have experienced an enormous push forward in the last years as computing power grew and the training of *deep neural networks* became a task with manageable time consumption [Sch15]. Hence, it's not surprising that people have proposed neural network driven approaches to computer composed music. The main advantage of those systems is their ability to capture repetitions of patterns over an arbitrary long period of time like Colombo et al. showed [CMS+16]. On the other hand, the main disadvantage of deep neural networks is their need for a high number of

[1] long short-term memories

examples (training data). Their results suggest that their model reproduces rhythmical structures well but is not able to capture high level melodic structure like phrase repetitions in rhythm or melody. Another example of a melody composing algorithm utilizing LSTMs among other approaches has been proposed by Coca et al. [CCZ13].

2.1.5. Generating Melodies with Markov Chains

Markov Chains are a statistical model introduced by the Russian mathematician A. A. Markov in the early 20^{th} century [Hay13]. They can be thought of as directed graphs where the vertices are events/states, edges are transitions and edge weights are the transition probabilities between events/states. In practical implementations Markov models are mostly represented with probability matrices. A main problem of Markov models used for algorithmic composition tasks is that the next state only depends on the current one. For melodies, this results in randomly wandering around pitch sequences. The results can be improved by using Markov models of higher order (m-th order). Instead of only depending on the current state, the current and the last $m-1$ states are taken into account for generating the next one.

There are two main ways to obtain the transition probabilities: either derive them by hand i.e. construct and tweak them until the results are satisfying, or learn them from a training set of melodies. The former approach was often taken by composers aiming for supporting their composition process while the latter one was focused on in more scientific contexts [FV13].

The training process for Markov models is easy. For all states, all transitions are counted and then divided by the total number or transitions. Here is a small example: our training data contains four different symbols “A”, “B”, “C”, and “D”. The training data consists of three documents that look like this: {“ABBA”, “ACDC”, “ACAB”}. For a second order Markov model, two blank symbols are prepended to all documents: {“␣␣ABBA”, “␣␣ACDC”, “␣␣ACAB”}. Now all following symbols for each bi-gram (that is not the last one of a document) are counted (see Table 1). Finally all rows are normalized to sum up to one. The result is called a transition matrix. It contains the probability of next states for a series of previous states.

It was already discovered in early years of the computer age [Moo72] that for lower order Markov models the resulting melodies wander aimlessly around in an unmusical way as mentioned above. For higher orders, one needs a big amount of training data or the results will be plain copies or at most recombinations of large musical fragments of the input. For intermediate orders the results are reasonable. One disadvantage, however, cannot be overcome by choosing a higher order Markov models: the composition process is only aware of a very local excerpt of it’s own composition. In this thesis we try to overcome this issue by modifying the melody composition process of the Markov model with other statistical concepts (see Section 4).

		next symbol			
		A	B	C	D
bi-gram	␣␣	3	0	0	0
	␣A	0	1	2	0
	AB	0	1	0	0
	BB	1	0	0	0
	AC	1	0	0	1
	CD	0	0	1	0
	CA	0	1	0	0

Table 1: An example of counting transitions to train a Markov model of second order. This transition matrix is calculated by normalizing each row to sum up to one.

2.2. Related Work

2.2.1. "Automatic Melody Generation" by Thayabaran Kathiresan

In 2015, Kathiresan proposed a work on composing melodies with first order Markov models trained with a single Melody [Kat15]. The Markov model was post processed with a CSP approach to include none or one of the four constraints: "Include a note ", "Include several notes ", "Use only a set of notes " and "Include a given pattern of notes ". Another approach using a hidden Markov model with categorical distributions was described as well. During evaluation, the first approach turned out to work better though. For this reason, we will focus on the formerly described approach throughout this work and call it KATH for shortness.

Kathiresan evaluated the work by conducting a listening test. The participants were asked to classify the results as "human composed ", "algorithm composed " or "not sure ". They were also asked to rate the pleasantness of the melodies on a scale from "excellent " over "good ", "fair " and "poor " to "bad ". His results are visualized in Table 2 and Figure 2.

		Predicted		
		Human Composed	Algorithm Composed	Not sure
Original	Human Composed	73 %	19 %	8 %
	Algorithm Composed	22.5 %	69 %	8.5 %

Table 2: Confusion matrix of the results of Kathiresan's listening test. The rows represent the true class of the melodies while the columns represent the class the participants assigned the the melodies presented.
Table extracted from Kathiresan's work [Kat15].

This thesis bases on the code and the idea but takes it further to overcome some of its main issues:

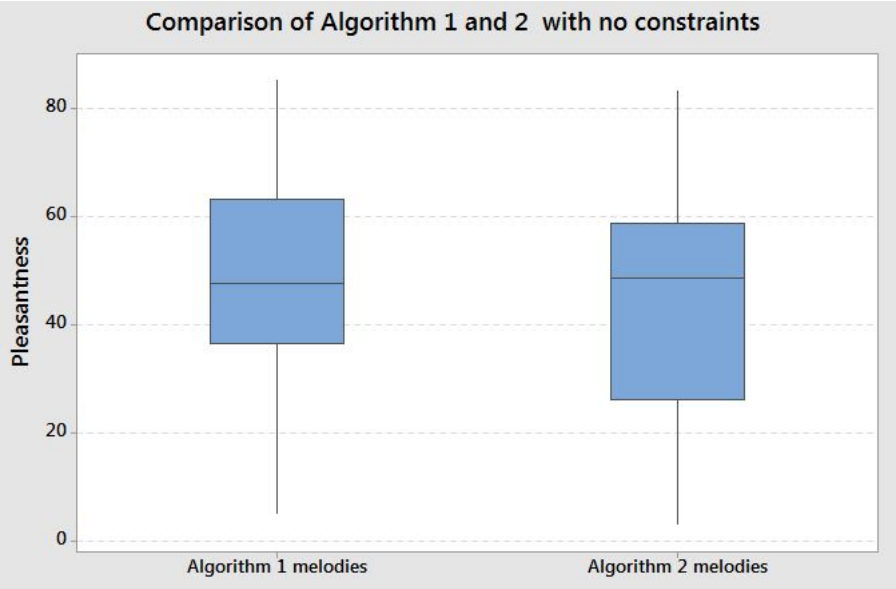

Figure 2: How pleasant participants of Kathiresan's listening test found the presented melodies. *Algorithm 1* is the one we used to compare our work to as it outperforms *Algorithm 2*.
Image extracted from Kathiresan's work [Kat15].

- Our main focus lies on compensating the lack of global composition awareness of Markov models as described in Section 2.1.5.
- We also try to restructure the composition process so that the length parameter to it can be given in bars instead of number of notes which makes less sense from a musical point of view.
- Another major point of improvement is that our algorithm is capable of learning from an arbitrary number of input melodies while Kathiresan's was restricted to one.
- We decide to deprecate the constraint feature in our code as it was implemented in a rather trivial manner with a lot of randomness and less systematic or musical considerations involved, leaving an improved implementation open as future work.

3. Musical Terms

3.1. Key and Mode

The *mode* of a musical piece defines a series of *intervals* (i.e. pitch distances) between the notes of a scale (i.e. consecutive pitches in a key—see below). Some examples of modes are *major* and *minor*. The corresponding scales can be found in Figure 3.

The key defines a the *root note* (i.e. where the scale starts) and the mode (see above) of a piece.

(a) *c major* scale (b) *c minor* scale

Figure 3: Scales of *c major* and *c minor* with tonal intervals.

3.2. Beats, Bars, and Time Signature

First, let's have a look at *note durations* (usually referred to as note values). There are *whole notes*[2] that can either be divided binary or ternary (using triplets) as shown in Figure 4. A time signature (e.g.) is defined by two integers. The lower one (always a power of 2) determines what note duration is considered as a *beat*, while the upper one determines how many beats are in a *bar*. The bar boundaries are indicated by horizontal lines in the staff, which is the system of line(s) the notes are written on.

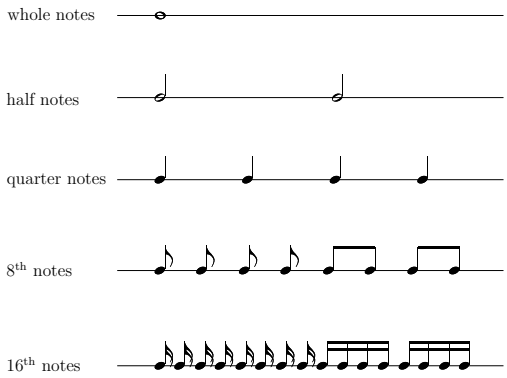

(a) Binary division of a *whole note*. Both representations of 8^{th} and 16^{th} notes are equivalent.

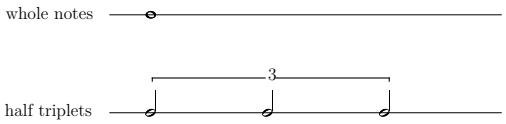

(b) A *whole note* can be divided into three *half triplets*. This applies to other note durations as well.

Figure 4: Note durations and their relations.

3.3. Count and Off-beat

In this work, *count* can either refer to a linguistic or a numerical representation of the position in a bar. Both representations share that successive integers, stating by 1, are assigned to the *beats* of each *bar*. While the numerical representation is then created by linear interpolation (see Figure 5a), the linguistic one is a little more complex. "and"s are inserted between the beats when the beat duration (see time signature in Section 3.2) is split in half. If split in half again, "a"s are inserted between the beats and the "and"s. If split into 3 (using triples), the new notes are called "trip" and "let" (see Figure 5b).

[2]Longer notes exist, but do not matter in this work.

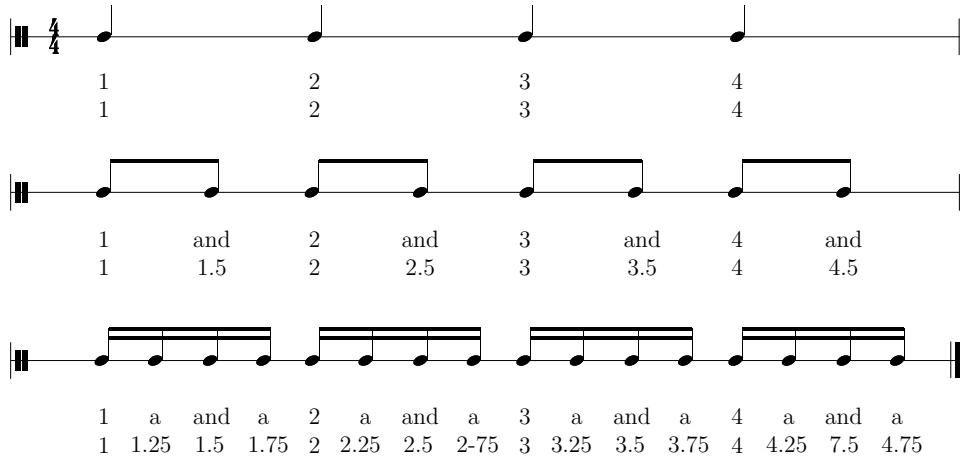

(a) Counting 16th notes with $\frac{4}{4}$ time signature.

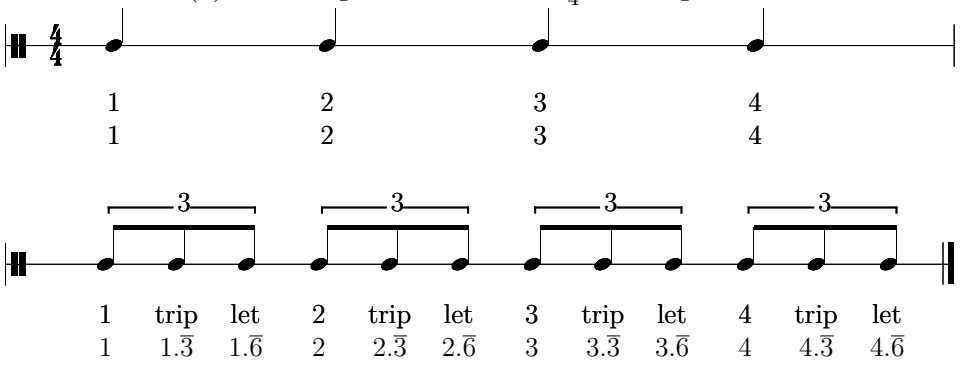

(b) Counting 8th triples with $\frac{4}{4}$ time signature.

Figure 5: Counting $\frac{4}{4}$ bars. The first lines of text are the linguistic way of counting, while the second lines are the numerical way.

When linguistically referring to a count a simplification is made for shortness. Instead of calling the 8th 16th note in Figure 5a (row 3) "1 a and a 2 a and a" all syllables before the last beat ("2") are dropped. Thus, it's called "2 a and a". Commonly syllables that aren't necessary for uniqueness are dropped as well, i.e. it's "2 and a". More generally, binary counts match the regular expression "`\d+( and)?( a)?`", while ternary counts match the regular expression "`\d+( trip( let)?)?`".

Though it usually refers to something different, we will use the term *off-beat* throughout the work as follows: The off-beat is the position in the bar, relative to the previous beat. Let n be the numerical representation of a count. Then the off-beat o is defined as n mod 1. The count "4 and a" (numerically 4.75) has an off-beat of .75.

4. Proposed Method

In the following, we will give an overview of the proposed method. The source code can be found on GitHub (`https://github.com/roba91/melody-composer`). We will explain the process more in detail in the next subsections. The proposed method is a hybrid approach combining Markov models (see Section 2.1.5) with a further statistical

extraction method aiming for compensating the lack of global context awareness of Markov models. Also, further modifications are made to give more context to the melody generation process. Our processing pipeline is shown in Figure 6.

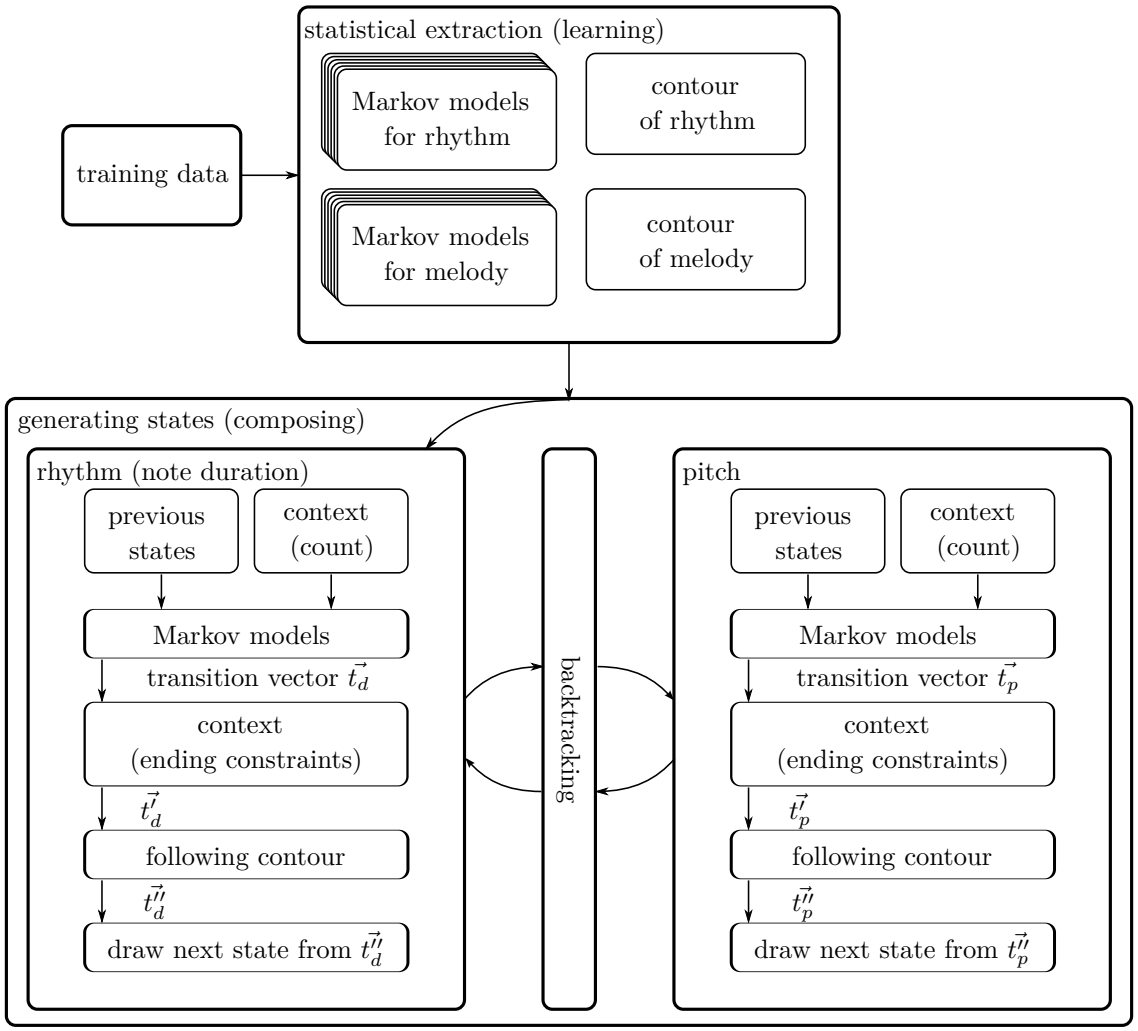

Figure 6: Workflow of the proposed method: An overview is given in Section 4. For details on the blocks see their respective sections. Training data: 4.1.1, Markov models: 2.1.5 and 4.1.2, Contour learning: 4.1.3, Ending constraints: 4.2.2, Following Contours and drawing next state: 4.2.3, Backtracking: 4.2.4.

The pipeline starts with the training dataset. In a pre-processing step, the melodies of the MTC-Dataset (see Section 4.1.1) are filtered (keeping only melodies in *major* mode with a $\frac{4}{4}$ time signature) and transposed[3] to the key of *C major*. This simplification was made to avoid having to handle different keys or modes and to increase the number of training melodies per key as all melodies are in *C major*. After that pre-processing the actual algorithm starts its work. It first reads in the melodies and splits them into

[3]Transposing a melody technically means to increase or decrease the pitch of all notes by a constant (e.g. one semi-tone).

phrases as indicated by phrase separator marks that came with the MTC-Dataset (see Section 4.1.1). The algorithm then extracts two statistical characteristics of rhythm and pitches for each melody:

- For note pitches and note durations each, a parametric Markov model of higher order is trained. Instead of using a transition matrix modeling a conventional Markov model we use a transition tensor. The added axis describes the current off-beat. It's defined within [0, 1). So for a note being played at count "*3 and a*" the value would be $3.75 \bmod 1 = 0.75$.

  The lower the order of a Markov model is, the less context sensitive will the results be. However, if the order is very high the results will be very close to the training data and the memory consumption increases exponentially with the order (see Section 4.1.2). The available RAM[4] determined our upper limit to Markov models of fourth order while a lower order makes less sense from a musical perspective as four notes often form a small musical pattern.

- For both pitches and note durations we extract a contour (see Section 4.1.3) for each phrase, which is basically the low-passed and x-normalized function of the feature (duration or pitch) over time. The contours are then clustered and one of the clusters is selected. A mean of all contours in it is calculated and used as contour to follow while generating melodies.

After that learning process, phrases can be generated. In contrast to the algorithm proposed by Kathiresan (see Section 2.2.1), where all pitches are generated first and the rhythm is generated afterwards, we generate a note duration, then the pitch for it and repeat that process until we are done. If the process reaches a point where the training data is insufficient (i.e. the combination of previous states couldn't be observed in the training data—at least not for the current count) we use backtracking to go back and truncate the search tree (see Section 4.2.4).

Now that a basic overview of the algorithm is given we will explain the details of it.

4.1. Learning

4.1.1. MTC-Dataset

Before learning anything, a dataset to learn from is needed. Hence, we will give an introduction to the MTC-Dataset that is used by the proposed algorithm.

As will explained in Section 4.1.3 we develop a method to extract melody and rhythm contours from our training data. Treating each song as a whole yields the problem that contours become very diverse and don't form nice clusters. However, considering phrase-wise contours works much better. Thus, we need a phrase annotated dataset.

[4]Random access memory: Temporal memory of a computer used for calculations.

The *Meertens Tune Collections* (MTC) [KBGW14] consists of several datasets of music pieces. The datasets vary in their representations and the meta data provided. The MTC-FS dataset is one of them. It consists of *digitally encoded* (many formats like MIDI[5], LilyPond[6], and PDF[7] are available) dutch vocal folk song *melodies* that are manually *phrase annotated*[8]. Thereby, it fulfills all our requirements.

Handling a dataset with all possible keys, modes and time signatures of melodies mixed together would clearly go beyond the temporal constraints of this work. Therefore, a pre-processing step that filters and normalizes the melodies to one key (*C major*) and one time signature ($\frac{4}{4}$) is introduced. It also removes polyphonic[9] input as our definition of melodies includes monophony.

The pre-processing tags the MIDI files with the difference of their key to *C major*[10]. A melody in *D major* e.g. would be tagged with the infix "*_m2*" (for minus two semitones) between the file name and the file ending. This infix is considered by our MIDI module when reading the files. It then transposes the melodies according to the file name infix.

The whole dataset has 4,120 songs with a total of 23,936 phrases (on average 5.81 phrases per song, average phrase length of 1.82 $\frac{4}{4}$-bars[11]) From those melodies we only use the ones in a *major* key with a $\frac{4}{4}$ time signature (find the distributions in Tables 3 and 4). After filtering the data, there were 378 songs with a total of 2147 phrases (on average 5.68 phrases per song, average phrase length of 1.92 $\frac{4}{4}$-bars) left.

major	minor	dorian	phrygian	lydian	mixolydian
3,879	161	69	4	2	5

Table 3: The mode distribution among the songs in the MTC-FS dataset.

4.1.2. Parametric Markov models

From a musical perspective it is intuitive that the probabilities of the note durations are different depending on their position in a musical phrase. Examples are phrase

[5]Musical Instrument Digital Interface: Its a multi-track (polyphinic) file format that represents musical events on a time scale. E.g. when and how long which note is to be played. See https://www.midi.org/.

[6]A file format that encodes sheet music in ASCII. See https://www.lilypond.org/.

[7]Portable Document Format: Its a cross-platform file format to encode documents. See https://www.adobe.com/devnet/pdf/.

[8]The phrase annotations are encoded as text events on a separate track in the MIDI files.

[9]Monophony means that at most one note is played at the same time. Polyphony makes no restrictions to number of simultaneous notes.

[10]The information about the key the melodies are in is taken from the LilyPond files that came with the MTC-Dataset. LilyPond files are designed to represent sheet music and therefore contain the key information. MIDI files in contrast, don't include information about the key, and guessing the key of a musical piece is error prone as it can be ambiguous.

[11]Evaluated on only 4,084 songs because 36 weren't usable as they were polyphonic, which our MIDI module isn't capable of handling.

k \ n	1	2	3	4	5	6	7	8	9	12
2	–	191	44	12	–	–	–	–	–	–
4	6	382	593	1,379	10	94	1	3	2	–
8	–	1	66	13	12	1,222	1	1	71	16

Table 4: The time signature ($\frac{n}{k}$) distribution among the melodies in the MTC-FS dataset.

ending rhythms and even more prominent the fact that it's very unlikely (at least for folk songs) that after two 8^{th} triples (on count "*1 trip let*") something else follows than another 8^8 triplet or similar (e.g. two 16^{th} triplets) because this would cause an offset by one 8^{th} triplet that wouldn't be resolved until the next 8^{th} triplet or similar. We will now present an approach to make Markov models aware of that fact.

Parametric Markov models can be thought of as several Markov models. Instead of a transition matrix a transition tensor of rank 3 is implemented. The first idea could be to use the added third dimension for the current position in a bar (i.e. the count) while composing. This causes the training process to need much more training data as transitions are separately observed for all counts in this case. From a musical perspective it can be argued that rhythm considerations are similar for same off-beats regardless of the count (one exception might be the count "*1*"). Thus, a reasonable simplification to this approach is to reduce the count to its off-beat. So, for a note being played at count "*3 and a*", the value would be $3.75 \bmod 1 = 0.75$. In other words, the third dimension does not represent the count (time relative to current bar), but the off-beat (time relative to current beat).

To get an idea of the memory consumption of our algorithm which is growing exponentially with the order chosen for the Markov model, here is an example. Let's assume we are looking at a parametric Markov model of fourth order for pitches and we have a total of 29 different pitches in our training set (which is the case for the MTC-FS dataset after filtering as described in Section 4.1.1). This would result in $\sum_{i=1}^{4} 29^i = 732,540$ rows each holding 29 probabilities (floats take up 32 bits = 4 bytes). So it uses $\sum_{i=2}^{5} 29^i \cdot 4\,\text{B} = 84,974,640\,\text{B} \approx 85.0\,\text{MB}$ per count. Assuming the training data contains a 32^{nd} as shortest note duration it means we have 8 of those transition matrices in our tensor. The storage needed is then $84,974,640\,\text{B} \cdot 8 = 679,797,120\,\text{B} \approx 679.8\,\text{MB}$. Repeating the same approximation with a Markov model of fifth order results in a memory usage of about 19.7 GB which isn't feasible any more for our computational infrastructure.

4.1.3. Contours

Markov models are known to not take a global context into account when used as a generative statistical model. To compensate for that, we will now introduce a contour concept. A contour can be thought of as the rough course of the melody (note pitches) or the rhythm (note durations) over time. The process to learn those contours for melody is very similar to the one for rhythm. Hence, we will focus on the melody contour in the following. The rhythm contour can be learned analogously.

Several approaches to melodic contours have been proposed. One of the earliest was described by Huron [Hur96]. In his work he only considers the relations between the first, last, and the mean pitch of a phrase. Thus, it is only capable to identify nine types of contours. It is not possible to use the model in a generative way, as it does not make use of absolute values, but ternary relations (higher, lower, same) of consecutive notes, while their duration (i.e. rhythm) is ignored. Schmuckler [Sch99] proposed a more suitable description of melodic contour. He uses the Fourier transform, which is an advantage in our use case, due to easier implementation of smoothing and comparison of contours, compared to Müllensiefen and Wiggens [MW11] who describe a polynomial approximation.

An illustration of the major steps of our method can be found in Figure 7. For each phrase we extracted a contour-feature vector like this:

1. Mirror the phrase. In other words: Prepend the reversed phrase to the original phrase. This avoids a possible wide jump from the last to the first note as the signal is assumed to be repetitive when calculating a discrete Fourier Transform.

2. Transform the phrase into a step function $m\colon [0,\, 1) \mapsto \{0, \ldots, 127\}$. Mapping a point of time to a MIDI pitch. (For the rhythm the function was defined as $r\colon [0,\, 1) \mapsto \mathbb{R}_{>0}$ where 1 was a quarter note, 0.5 a quaver and so on.)

3. Center the phrase to be around 0. This step helps to learn contours, regardless of their register. This step is omitted for rhythm. Several methods for transposing were tested: Transposing so that the first note (that isn't a rest) is 0, transposing so that the mean of the melody is 0 and transposing by the mean of all notes (ignoring their duration). The latter method lead to the best clustering behavior for our dataset as indicated by first experiments. A final evaluation of the results is left as future work.

4. Sample the function in the time axis.

5. Interpolate rests. Rests don't have a pitch but we needed m to be a function $(\forall t \in [0,\, 1)\colon \exists m(t))$ for the next steps. This interpolation step is only necessary for the melody, since rests define a rhythm as well. For rests at the beginning of the phrase the first note with a ,pitch is expanded to the beginning of the phrase.

For those at the end the last note is expanded. For rests between pitches the interpolation is done linearly.

6. Apply the Fast Fourier Transform (FFT) to the function. Then set all frequencies above a threshold to zero. First experiments indicate that keeping the lowest 6 frequencies is a good choice. An evaluation of the impact of different choices is left for future work (see Section 6). This acts as a low-pass filter (smoothing of the curve). It is important to beware of the energy-loss that such a operation causes. So we multiplied the result with the ration of energy lost to not lose tonal or rhythmic range (i.e. the width of the function when transformed back into the time domain). For implementation details see *MelodyContour.low_pass* in *contour.py*.

These steps result in a feature vector $\vec{x} \in \mathbb{C}^6$ for each phrase. Unsupervised learning techniques can then be applied to find a representative contour for composition:

7. Clustering of the vectors. You can think of the clusters as families of contours. Using the Ward variance minimization algorithm[WJ63] yields the best looking results in our use-case (see Figure 8). Ward's algorithm is a hierarchical clustering algorithm that, in each step, joins the cluster pair that causes the least increase of in-cluster variance. Different algorithms like K-Means and hierarchical clustering with all distance measures provided by SciPy[12] including min-, max-, and average-distance result in more diverse clusters. For this work a number of up to 17 clusters turned out to be a good choice as indicated by small experiments. However, in future work a more refined method of choosing the clustering threshold depending on the presented training data is possible (see Section 6).

8. Replace the cluster members with their not transposed matches (see Step 3).

9. After clustering one of the resulting clusters is chosen. A first trivial way to do so is to select the largest cluster. This however, often results in a cluster with a mean curve (see Step 10) that has little variance. A more promising approach is to use a quality measure incorporating size and variance of the mean curve to choose a cluster. For each cluster C_i let

[12]For more detail see the SciPy-Documentation: http://docs.scipy.org/doc/scipy/reference/generated/scipy.cluster.hierarchy.linkage.html#scipy.cluster.hierarchy.linkage

$$r_{C_i}(t) = \mathrm{FFT}^{-1}\left(\frac{1}{|C_i|}\sum_{\vec{x_i}\in C_i}\vec{x_i}\right) \quad \text{be the mean contour,}$$

$$s(C_i) = \frac{\|C_i\|}{\max_j |C_j|} \quad \text{the normalized size,}$$

$$mean(r_{C_i}) = \int_0^1 r_{C_i}(t)\,dt \quad \text{the mean of } r_{C_i} \text{ and}$$

$$w(C_i) = \int_0^1 |mean(r_{C_i}) - r_{C_i}(t)|\,dt \quad \text{the width of the cluster.}$$

The quality measure is then defined as $q(C_i) = \frac{s(C_i)}{\max_j s(C_j)} + \frac{w(C_i)^{\frac{1}{\gamma}}}{\max_k w(C_k)}$ where γ is the weighting parameter, determining the ratio of importance between a cluster's size and the width of its mean contour. $\gamma = 3$ is chosen due to preliminary experiments. An exact evaluation of the best choice for γ is left as future work (see Section 6). We then select the cluster C_i by $\arg\min_i q(C_i)$.

10. The learned contour is then defined as the inverse FFT of the mean contour of the selected Cluster after a mapping to invert the mirroring (see Step 1): $contour\colon [0,\,1) \mapsto \{0,\,...,\,127\}, t \mapsto r_{C_i}\left(\frac{1+t}{2}\right)$

4.2. Phrase Composition

4.2.1. Joining Pitch and Rhythm Generation

Following the human way of composing, we change KATH's method of composition (see Section 2.2.1). KATH is only capable of generating a melody with a given number of notes but not a given number n of bars or beats, it composes melodies by first generating n pitches and then n note durations. By changing this strategy to iteratively generate a single duration and then a single pitch, better control over the duration of the melody as well as more context for the composition process is gained. The new algorithm has the information of where in a bar or the piece it is all the time.

With this gain of control over the duration of the melody – in terms of not needing to know the number of notes we want to use for the composition before knowing the rhythm or the melody – it is a straight forward modification to the algorithm to make it work with a number of $\frac{4}{4}$-bars as parameter instead of number of notes.

The knowledge of *where* or better *when* the algorithm is within the composition is used to select a Markov model according to the current off-beat. As explained in Section 4.1.2 the learning process incorporates training parametric Markov models. The off-beat is used as parameter for the models. The off-beat is calculated as c mod 1, where c is the count in the current bar. The knowledge of the position relative to the whole composition is also used to follow the learned contours (see Section 4.1.3) as will be described in Section 4.2.3.

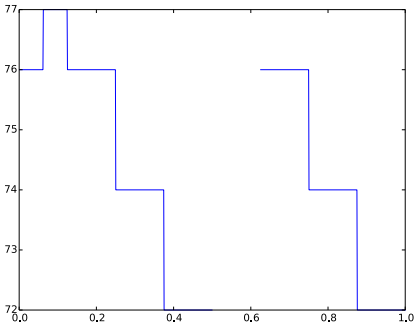

(a) A phrase's melody.

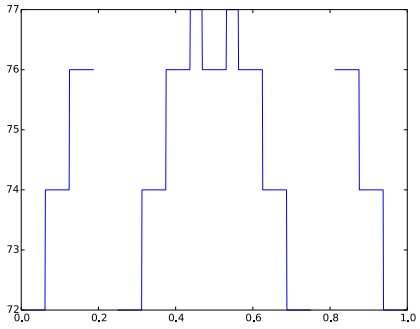

(b) Mirroring the phrase (Step 1).

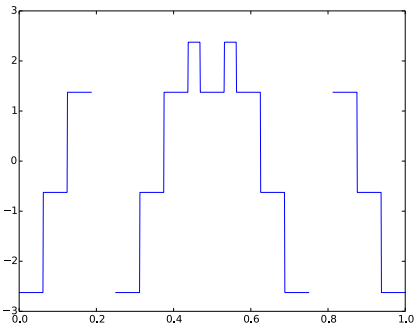

(c) Transposing the phrase (Step 3).

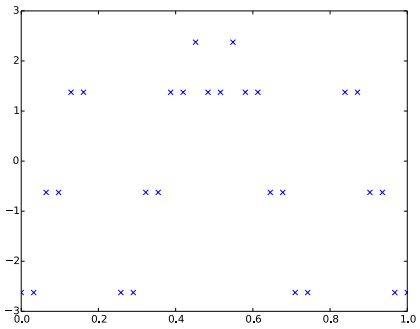

(d) Time discretize the phrase (Step 4).

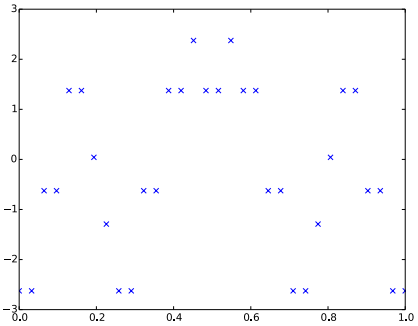

(e) Interpolate the rests (Step 5).

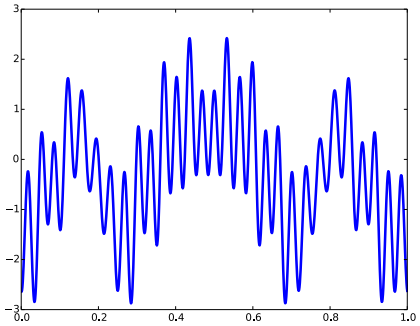

(f) Fourier transform the phrase† (Step 6)

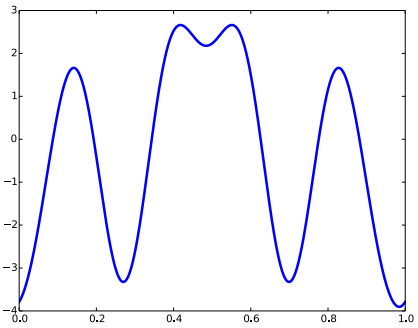

(g) Apply low pass filter† (Step 6).

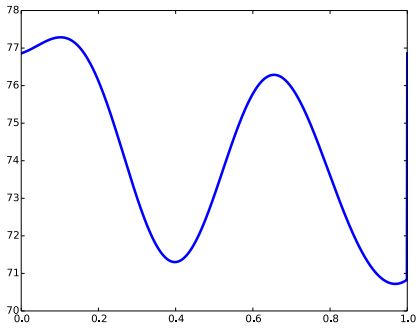

(h) Unmirror the result (Step 10).

Figure 7: A minimal example of learning a melody contour. For simplicity we only use one melody. Because of this, the whole clustering part (steps 7–9) is skipped. †While these steps are actually performed in the frequency domain (after applying the FFT) we chose to plot them in the time domain (applying FFT^{-1}) for better illustration here.

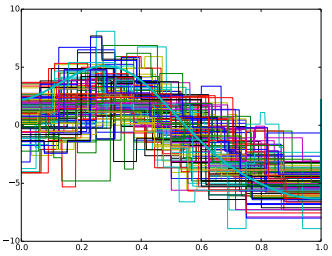

(a) $|C_1| = 115$; $q(C_1) \approx 1.26$

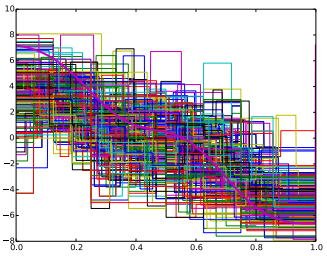

(b) $|C_2| = 235$; $q(C_2) \approx 1.68$

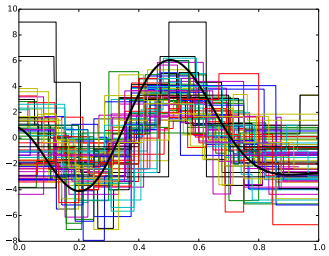

(c) $|C_3| = 69$; $q(C_3) \approx 1.01$

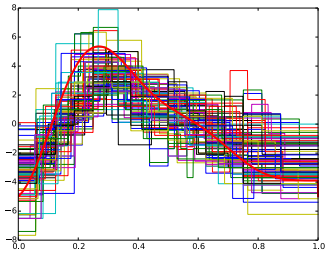

(d) $|C_4| = 93$; $q(C_4) \approx 1.10$

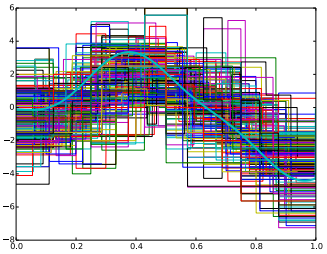

(e) $|C_5| = 283$; $q(C_5) \approx 1.68$

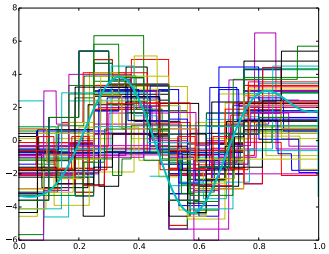

(f) $|C_6| = 73$; $q(C_6) \approx 0.99$

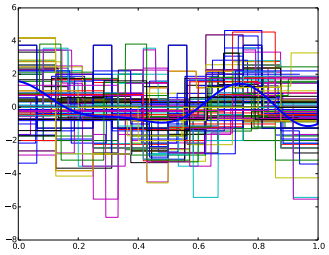

(g) $|C_7| = 126$; $q(C_7) \approx 0.94$

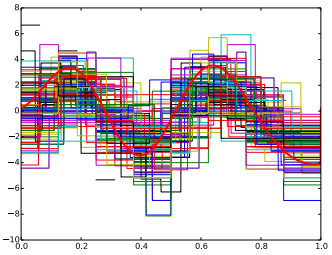

(h) $|C_8| = 128$; $q(C_8) \approx 1.17$

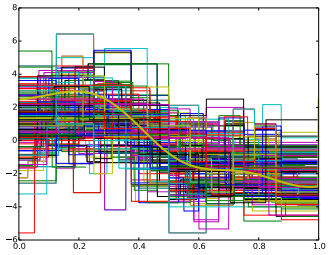

(i) $|C_9| = 236$; $q(C_9) \approx 1.53$

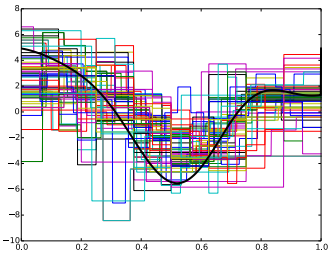

(j) $|C_{10}| = 62$; $q(C_{10}) \approx 0.98$

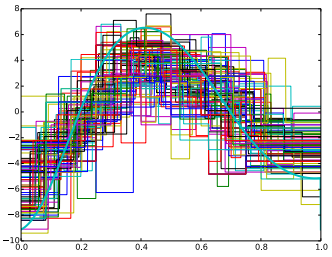

(k) $|C_{11}| = 94$; $q(C_{11}) \approx 1.22$

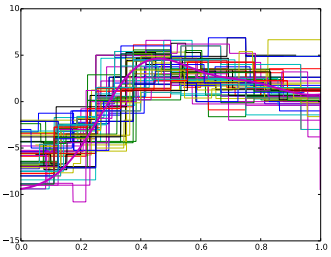

(l) $|C_{12}| = 60$; $q(C_{12}) \approx 1.05$

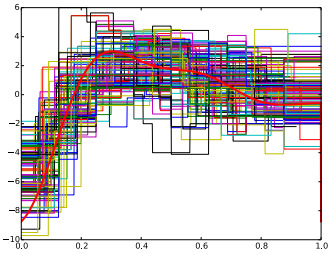

(m) $|C_{13}| = 135$; $q(C_{13}) \approx 1.16$

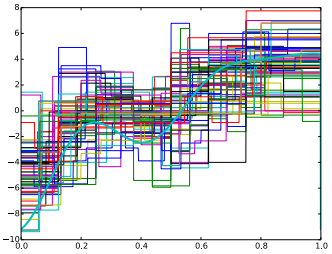

(n) $|C_{14}| = 73$; $q(C_{14}) \approx 1.06$

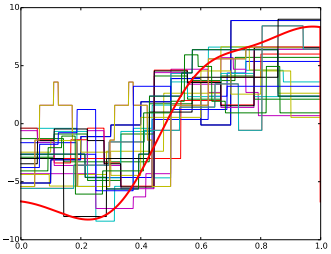

(o) $|C_{15}| = 23$; $q(C_{15}) \approx 1.08$

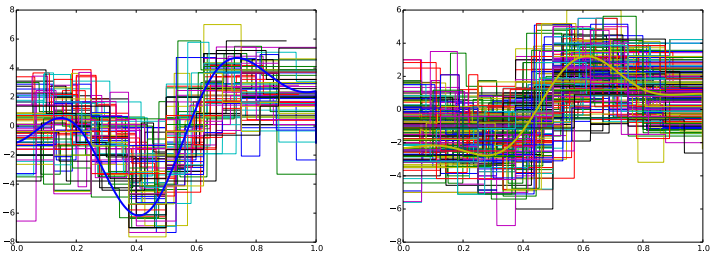

(p) $|C_{16}| = 77$; $q(C_{16}) \approx 1.04$ (q) $|C_{17}| = 229$; $q(C_{17}) \approx 1.49$

Figure 8: All phrases of the MTC-FS data set clustered by their melody's contour. $|C_i|$ is the number of phrases in cluster i while $q(C_i)$ is its quality measure as defined in Section 4.1.3 Step 9. In this example, cluster C_5 will be selected as $\arg\max_i q(C_i) = 5$.

After selecting the appropriate Markov model according to the current off-beat it defines a transition vector $\vec{t}$ for each state (or sequence of states if using higher order Markov models). $\vec{t}$ contains the probability distribution[13] of transition for each possible state. We call the transition vector which depends on the off-beat and the previous sequence of states $\vec{t_p}$ if we refer to pitches as states, $\vec{t_d}$ if we refer to durations as states, and $\vec{t}$ if we talk about both kinds of transition vectors.

Let $\vec{s}$ be the vector of all sates of the Markov model. Note that states are numerical in our case (durations as well as pitches)[14]. We will use $\vec{s_p}$ when talking about pitches as states and $\vec{s_d}$ for durations as states.

4.2.2. Ending Constraints

Assume $\left\|\vec{t}\right\| = 1$ as all cases where $\left\|\vec{t}\right\| = 0$ are caught before as will be described in Section 4.2.4. In the next step, we apply ending constraints to $\vec{t_p}$ to obtain $\vec{t'_p}$ and to $\vec{t_d}$ to obtain $\vec{t'_d}$.

- The duration transition vector $\vec{t_d}$ is manipulated in a way to make sure that no longer duration is generated than time (counts) is left in the composition. Let c be the current count in the composition and c_{total} the wanted length of our composition in counts. We now define a duration manipulation vector $\vec{m_d} \in \mathbb{R}^{dim(\vec{s_d})}$ element-wise as $\vec{m_{d_i}} = \begin{cases} 1 & \vec{s_{d_i}} \leq c_{total} - c \\ 0 & \text{else} \end{cases}$. We then calculate $\vec{t'_d} = \left\langle \vec{t_d} \odot \vec{m_d} \right\rangle$.
- The pitch vector is manipulated in a way, that the last pitch in a melody is one of the tonic triad. As we are working only in *C major* this is *c*, *e* and *g* in any octave. Since the rhythm is generated first, the algorithm can easily determine

[13] As it's a distribution $\left\|\vec{t}\right\| = 1$. It may happen though, that $\left\|\vec{t}\right\| = 0$. This is the case, when the training data does not contain any transition from the corresponding (sequence of) state(s).

[14] The MIDI format encodes pitches with integers. We make use of that here. Durations are normalized to quarter notes. So a quarter note has the value 1 while a eight note has the value 0.5 etc.

which pitch is the last one. This is calculated as

$$\vec{t'_p} = \begin{cases} \left\langle \vec{t_p} \odot \vec{m_p} \right\rangle & \text{if last pitch} \\ \vec{t_p} & \text{else} \end{cases}, \text{ with } \vec{m_{p_i}} = \begin{cases} 1 & \vec{s_{p_i}} \in \{\text{c, e, g}\} \\ 0 & \text{else} \end{cases} \in \mathbb{R}^{\dim(\vec{s_p})}$$

4.2.3. Following Contours

Now that we have two contour functions over time that describe the course of the composition, one for its melody and one for its rhythm (see Section 4.1.3). We have to manipulate the trained Markov models so that their output follows the respective contour. The basic idea here is to manipulate the transition probabilities with a Gaussian filter and let its center follow the learned contours. This will make the melody and the rhythm loosely follow the contours as well.

We therefore define a filter vector $\vec{f} \in \mathbb{R}^{|\vec{s}|}$ element-wise as $\vec{f_i} = \varphi_{\mu,\sigma^2}(\vec{s_i})$ where μ is the value of the contour at the current time in the composition and σ^2 a tweaking parameter. $\sigma^2 = 4$ for melody contours (which means a standard deviation is four semi-tones) and $\sigma^2 = 0.33$ for the rhythm contours (which means a standard deviation of an eighth triplet) are chosen as preliminary experiments showed they work well. A full evaluation of the effects and the best choices is left as future work (see Section 6).

Let $\vec{t'}$ be the transition vector as described in Section 4.2.2. We then calculate the filtered transition vector $\vec{t'} = \left\langle \vec{t} \odot \vec{f} \right\rangle$ and use this instead to draw the next state[15].

4.2.4. Backtracking

As the algorithm composes a melody, it might find itself in a situation where the previous states do not appear in the training data. In that case, $\vec{t} = \vec{0}$. All the manipulations in Sections 4.2.2 and 4.1.3 would fail as $\left\langle \vec{0} \right\rangle$ implies a division by zero. Skipping the manipulations wouldn't help either as there is no information on which state could or should be generated next. A way around that problem is to use backtracking. Since the parametric Markov model's transition probabilities depend on the count, it might well happen that the problem of insufficient training data can be resolved by just drawing another duration state while keeping the pitch. Therefore, our algorithm sticks to one pitch after failing and will not change that pitch until all possible durations for it have failed. This can be quite memory consuming if the training data is small. For every leaf in the search tree (which corresponds to an instance of insufficient data) the algorithm stores the current pitch (the one we decided to try all durations for) and all durations tried yet for each pitch tried yet in its parent node. For the exact implementation, see the code at *melody_generation.py:MelodyTreeTruncate*.

[15] States are, like in the original work [Kat15], drawn by choosing a random number $r \in [0, 1]$, calculating the accumulative sum $\vec{s} = accsum(\vec{t'})$ and returning $\vec{s_i}$ with the smallest i, $r \leq \vec{s_i}$.

5. Results

We now want to evaluate the results of the algorithm proposed in Section 4. The results will be evaluated on the basis of a human benchmark incorporating aesthetics and ambiguity error. Therefore, we will first create a baseline from KATH and the training data (see Section 5.1). We will then explain the online survey set up (see Section 5.2), and finally, we will present and interpret the results of the survey (see Sections 5.3 and 5.4).

5.1. Generating the Baseline

The algorithm proposed by Kathiresan [Kat15] is only capable of learning from one MIDI file and composes a melody with a given number of notes. To be comparable to our algorithm a set of ten melodies was created with our algorithm. It was run five times with only one melody as training data, composing two melodies with a length of four $\frac{4}{4}$-bars each time. For Kathiresan's algorithm a thin wrapper is needed to force it to compose melodies with four $\frac{4}{4}$-bars. The wrapper reads the training data and calculates from the longest and the shortest note in the data the minimal and maximal number of notes that fit into four $\frac{4}{4}$-bars. In a loop a random number n in that range is drawn. Kathiresan's algorithm is then called to compose a melody with n notes. This is repeated until the composed melody has a length of four $\frac{4}{4}$-bars. The same training melodies used for our algorithm are used for Kathiresan's as well to achieve comparability of the results.

5.2. Conducted Listening Test

The conducted listening test was set up as an online survey. This is possible as there are no requirements to be met by the participants except for the ability to hear. Nor are requirements to be met by the subject's environment expect the ability to hear the presented stimuli. To guarantee that a test stimulus was presented to all subjects, before they started the survey. They were asked if they could hear it. If not, we excluded them from our study.

One of the main advantages of an online survey for our case is the fact that it is much easier to acquire a higher number of participants than with a face-to-face test. Knowing that online surveys will most likely not result in a representative sample of the population [Wri05] we prefer having a big sample over a small but representative one. To avoid costs, the server software is implemented by ourselves and hosted on already existing infrastructure.

To prevent over-tiring of the test subjects, the test time was limited. Therefore, the subjects were randomly split into five groups. Each subject was first asked to answer a few personal questions. A melody was then presented and the subjects were asked

three questions about it: 1) if the melody was known to the them, 2) if the subjects think the melody is human or computer composed, and 3) how much the subjects like the melody). This was repeated for twelve different melodies. The twelve melodies presented to each subject were different for each group and in random order for each subject, but the distribution of types were the same for all groups. See Table 5. The exact questions asked can be found in Appendix 6.

type	qty.	source (type)	parameters	comment
1	2	Randomly picked from MTC dataset	Phrases of length four $\frac{4}{4}$-bars	
2	2	Manually composed by us for this test	Phrases of length four $\frac{4}{4}$-bars in a similar style to phrases of the MTC dataset	
3	2	Generated from Kathiresan's work [Kat15]	Length forced to be four $\frac{4}{4}$-bars	see Section 5.1
4	2	Generated by our algorithm	Learn from one mtc-melody only	To be comparable to Kathiresan's work
5	1	—"—	from 10 melodies	
6	1	—"—	—"— 20 —"—	
7	1	—"—	—"— 50 —"—	
8	1	—"—	—"— all (378) —"—	

Table 5: This table shows the distribution of melody types presented to each subject regardless of the group assigned.
All melodies composed by our algorithm implicitly share the following parameters: length of four $\frac{4}{4}$-bars and Markov model order four.
Please take note of the difference between the terms "melody" and "phrase", especially when talking about the MTC dataset. A melody can be composed of multiple phrases.

5.3. Results of the Listening Test

The online survey was open for about 6 weeks. A total of 248 subjects participated. In the following the answers are visualized.

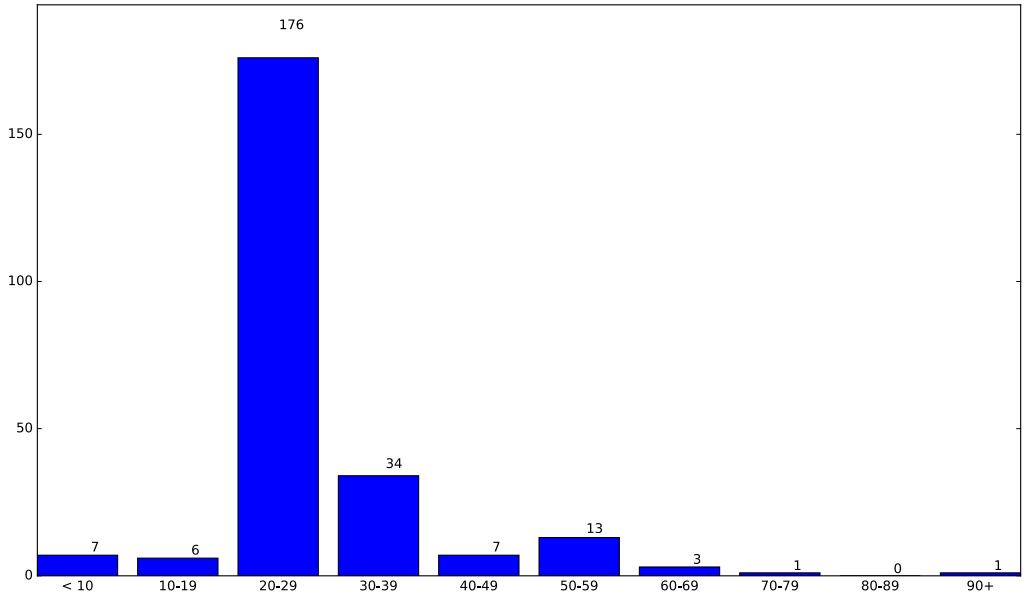

Figure 9: Age of participants.

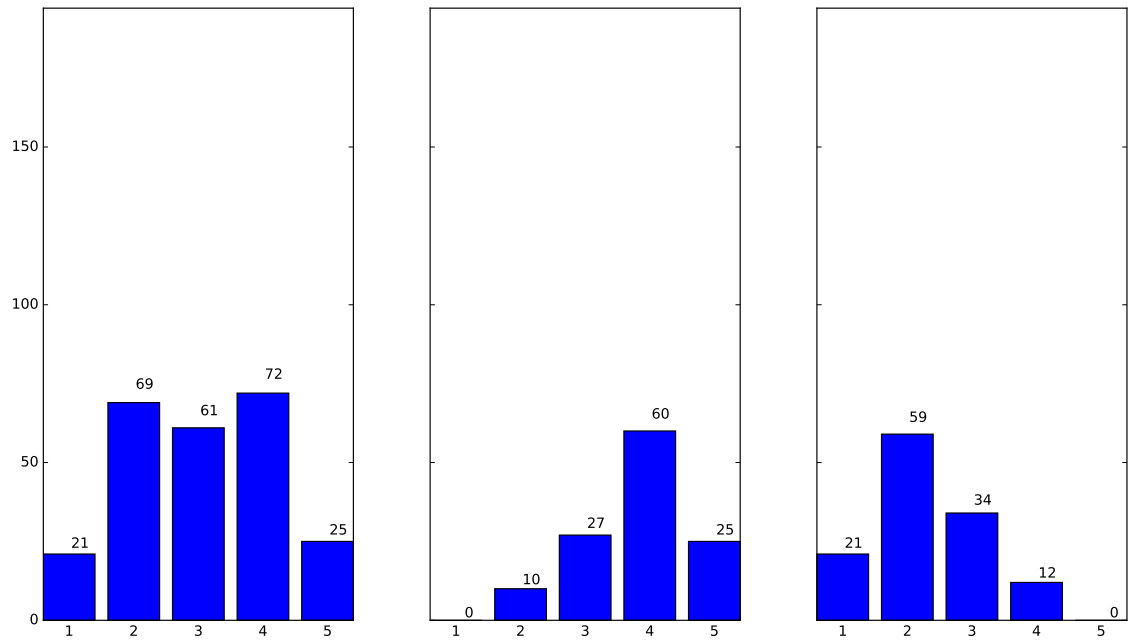

(a) all
$\mu = 3.04,\ \sigma^2 = 1.14$

(b) playing an instrument
$\mu = 3.82,\ \sigma^2 = 0.85$

(c) playing no instrument
$\mu = 2.29,\ \sigma^2 = 0.86$

Figure 10: Self-estimated musicality of the subjects. The subjects were asked to self-estimate their musicality on this scale:
1 – “not at all”, 2 – “a bit”, 3 – “moderate”, 4 – “good”, 5 – “excellent”

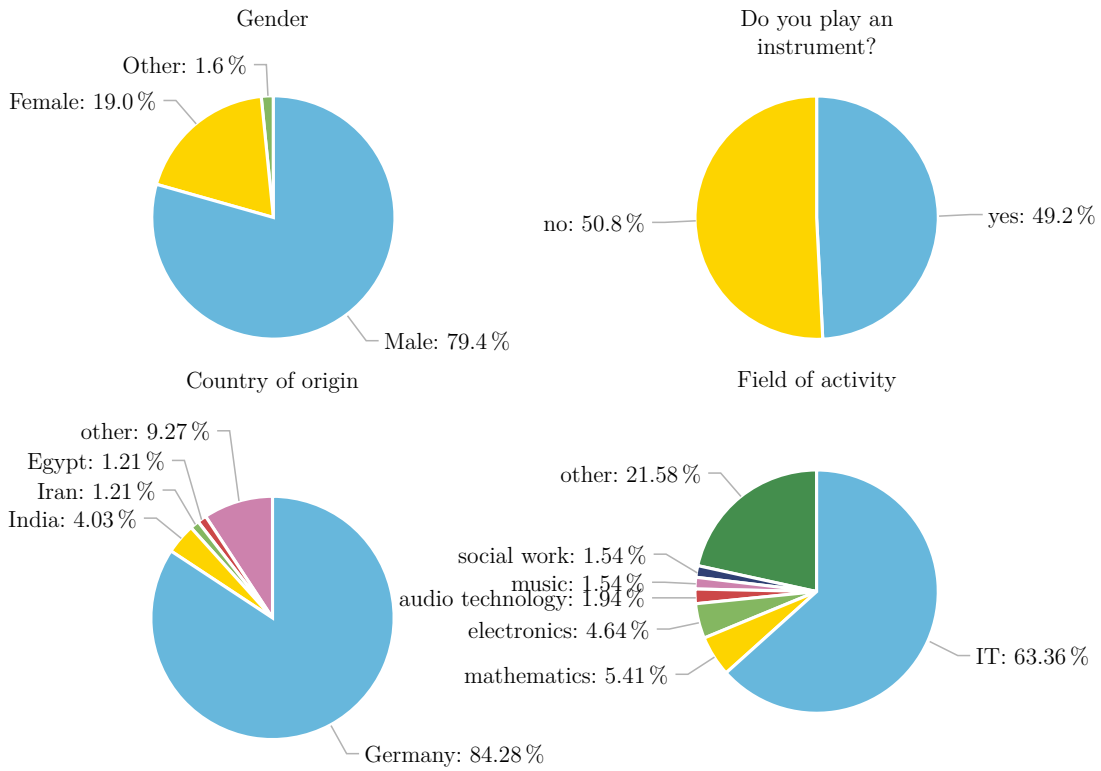

Figure 11: Distribution of gender, instrument playing, "field of activity" and origin among subjects.

ground truth		**answers**						
type	human	Σ	k	h	$k \wedge h$	$\neg k \wedge h$	$k \wedge \neg h$	$\neg k \wedge \neg h$
1	yes	496	76	343	74	269	2	151
2	yes	496	52	348	46	302	6	142
3	no	496	11	97	4	93	7	392
4	no	496	67	261	60	201	7	228
5	no	248	20	151	20	131	0	97
6	no	248	11	123	7	116	4	121
7	no	248	8	124	6	118	2	122
8	no	248	16	133	14	119	2	113

Table 6: Σ is the number of responses for the given type of melody. k is the number of those responses that stated to know the melody. h is the number that stated the melody to be composed by a human.

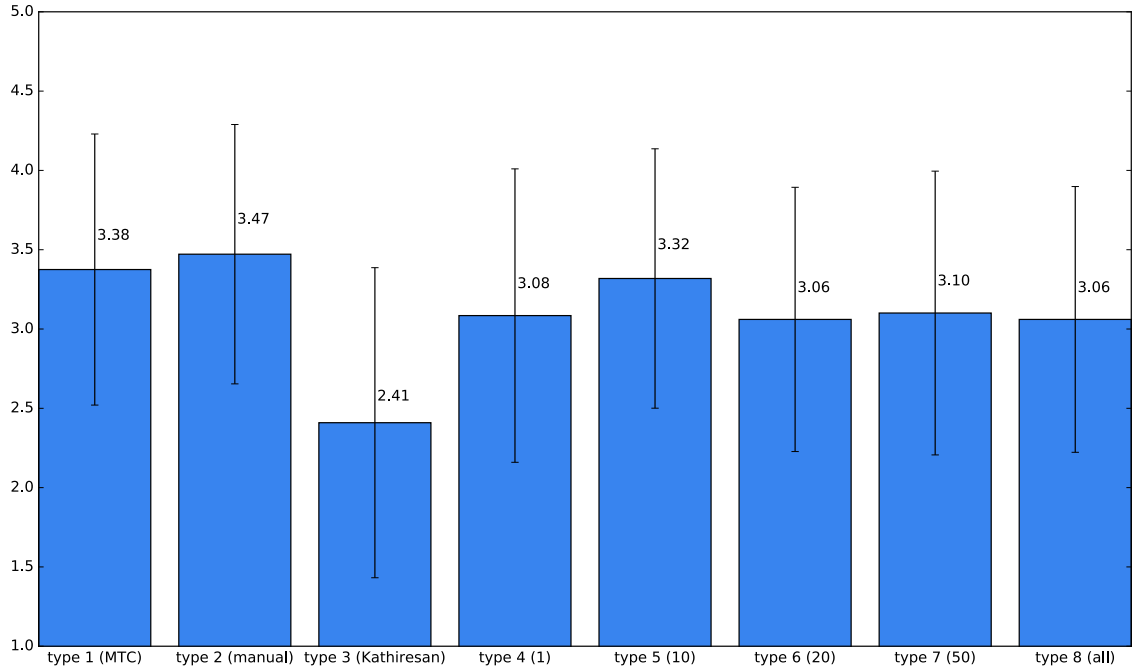

Figure 12: Subjects had to rate how much they like every melody on a scale from 1 to 5 ("strongly dislike","dislike","neither like nor dislike","like","strongly like").

5.4. Interpretation of the Results

As mentioned in Section 5.2 it was expected that the probe of people would not be representative. Viewing the age and gender distribution (Figures 11 and 9) compared to the German 'Zensus 2011' [Sta11] proves that. Our subjects are younger and there are more males among them than the German average shows. This is probably due to the fact that the listening test was mainly advertised at the local university with a stress on the computer science department (clearly to see in Figure 11 "field of activity"). However, the number of subjects playing a music instrument is more than 30 % higher than the German average of 17 % [Mus14]. It can be assumed that subjects with musical experience are better in detecting non-human composed melodies. Thus, one should keep in mind that our subjects were probably more skeptical towards the presented melodies when classifying them as human or computer composed. Another aspect becomes visible when analyzing the *age* and the *field of activity* responses: When opening the access to the internet without restrictions, the responses will end up having some noise. It does not seem plausible that we reached seven people younger than ten years old, nor people older than 90 that are able to operate a computer. This becomes even more clear when "sdfgh" is answered when asking for the *field of activity*.

Looking at Table 7 we first want to see how well our subjects were able to distinguish human versus computer composed melodies. The row $\frac{\neg k \wedge h}{\neg k}$ of Table 7 shows how often a phrase was classified as "human composed", excluding all cases where the subject knew the phrase. These cases are excluded because the subjects can deduce from knowing the

ground truth		answers		
type	human	$\frac{k}{\Sigma}$	$\frac{h}{\Sigma}$	$\frac{\neg k \wedge h}{\neg k}$
1	yes	15.32 %	69.15 %	64.05 %
2	yes	10.48 %	70.16 %	68.02 %
3	no	2.22 %	19.56 %	19.18 %
4	no	13.51 %	52.62 %	46.85 %
5	no	8.06 %	60.89 %	57.46 %
6	no	4.44 %	49.60 %	48.95 %
7	no	3.23 %	50.00 %	49.17 %
8	no	6.45 %	53.63 %	51.29 %

Table 7: Percental values relevant for evaluation of the results. They were calculated from Table 6.

phrase that it must be human composed. However, we only want to look at deduction from the phrase itself. The results for types 1 to 3 match the results Kathiresan found (see Table 2). While KATH could only reach about 20 %, our proposed algorithm reached about 50 % of subjects thinking the phrase was human composed. This is as good as random guessing. So the subjects were not able to decide if the presented phrases were human or computer composed. Comparing this number with those of types one and two, it is interesting to see that the participants did not do much better for the human composed phrases.

The answers visualized in Figure 12 yield similar results. In fact the data is nearly perfectly linear correlated as shown in Figure 13.

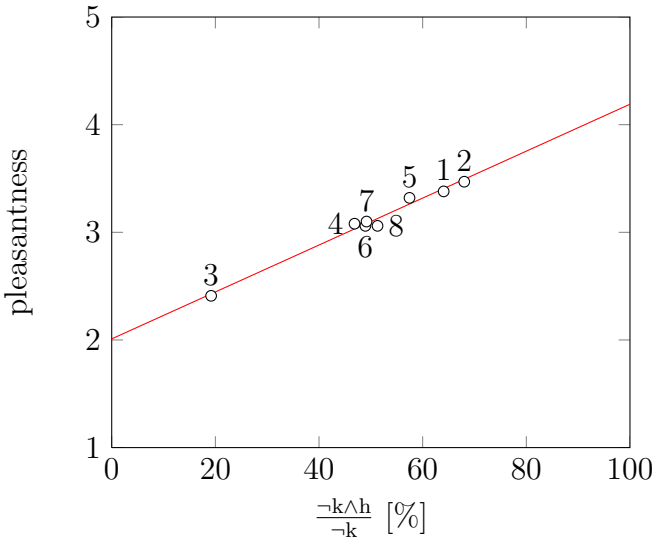

Figure 13: Showing the correlation between how the subjects liked a set of phrases and how often they thought the melody was human composed. The red line shows the linear regression of the data points. For the null hypothesis that the error is 0 the two-sided p-value is $1.5 \cdot 10^{-6}$; the standard error is 0.12.

Finally, a side note on the high number of subjects stating to know the phrases of type four presented to them, even though they were composed by our algorithm: Since this type of phrases has only one melody as reference (training data), it isn't surprising that the results are very close to the input. Looking at Appendix 6, one can see that this is true indeed but they are not simple replicates of the input. This is theoretically possible as this is not checked explicitly by the proposed algorithm.

6. Conclusion and Future Work

In this work, we have shown that our system produces reasonable results, significantly outperforming the given baselines. Even though Markov models alone are seen as no proper method for algorithmic composition, we successfully showed that when combined with further methods they can yield much better results in terms of being closer to human composed melodies. This can be seen when comparing our results with the ones of Kathiresan [Kat15], whose basic algorithm solely relies on Markov models. Apart from the previous works, our algorithm outperforms a random guessing baseline, meaning that humans are not able to clearly distinguish its compositions from humans anymore.

Other aspects of melody composition are still open and could improve the results. Due to time restrictions, we had to leave those other aspects as future work. The most promising areas of future study could be the following:

- Generate more but shorter phrases to compose a livelier and longer melody. One could incorporate concepts like variation, repetition and similar in this process to imitate the human composition processes.

- A method for finding proper phrase endings in terms of rhythm. As is, our algorithm might well end a phrase with a very atypical rhythmical pattern, because the transition probability vector of the duration Markov model is not modified except for the last note to make sure the composition does not get too long. A possible solution to this could be to compose the melody from beginning to end and from end to beginning (with a separately trained Markov model) simultaneously and meet somewhere in the middle of the composition.

- Enrich musical diversity by supporting different keys, modes (major, minor, etc.) and time-signatures.

- The contour clustering uses a fixed number of clusters. However, in future work a more refined method of choosing the clustering threshold depending on the presented training data could improve the algorithm.

- Utilize a similarity function to avoid plain rehashing or replicating of the input data.
- Several tweaking parameters were chosen by small experiments because the resources for an intense study to find the best choices were missing. Namely the low-pass filter threshold in Section 4.1.3, Step 6, the weighting parameter γ in Section 4.1.3, Step 9, and the standard deviation values in section 4.2.3 that determine how accurately the algorithm follows the learned contours.

The interested reader can find the source code of the proposed method on GitHub (`https://github.com/roba91/melody-composer`).

Bibliography

[CCZ13] Andrés E Coca, Débora C Corrêa, and Liang Zhao. Computer-aided music composition with lstm neural network and chaotic inspiration. In *Neural Networks (IJCNN), The 2013 International Joint Conference on*, pages 1–7. IEEE, 2013.

[CMS+16] Florian Colombo, Samuel P Muscinelli, Alexander Seeholzer, Johanni Brea, and Wulfram Gerstner. Algorithmic composition of melodies with deep recurrent neural networks. *arXiv preprint arXiv:1606.07251*, 2016.

[DHS12] Richard O Duda, Peter E Hart, and David G Stork. *Pattern classification*. John Wiley & Sons, 2012.

[ELPS+07] D Espı, PJ Ponce de León, Carlos Pérez-Sancho, David Rizo, José Manuel Inesta, F Moreno-Seco, and Antonio Pertusa. A cooperative approach to style-oriented music composition. In *Proc. of the Int. Workshop on Artificial Intelligence and Music, MUSIC-AI*, pages 25–36, 2007.

[FGW13] Sebastian Flossmann, Maarten Grachten, and Gerhard Widmer. Expressive performance rendering with probabilistic models. In *Guide to Computing for Expressive Music Performance*, pages 75–98. Springer, 2013.

[FV13] Jose David Fernández and Francisco Vico. Ai methods in algorithmic composition: A comprehensive survey. *Journal of Artificial Intelligence Research*, 48:513–582, 11 2013.

[GTK09] Jon Gillick, Kevin Tang, and Robert M Keller. Learning jazz grammars. In *Proceedings of the Sound and Music Computing Conference*, pages 125–130, 2009.

[Hay13] Brian Hayes. First links in the markov chain. *American Scientist*, 101:92–97, March—April 2013.

[HS97] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.

[Hur96] David Huron. The melodic arch in western folksongs. *Computing in Musicology*, 10:3–23, 1996.

[Jen11] Johannes Høydahl Jensen. Evolutionary music composition: A quantitative approach. Master's thesis, Norwegian University of Science and Technology, Institutt for datateknikk og informasjonsvitenskap, 2011.

[Kat15] Thayabaran Kathiresan. Automatic melody generation. Master's thesis, KTH Royal Institute of Technology, 2015.

[KBGW14] Peter van Kranenburg, Martine de Bruin, Louis P. Grijp, and Frans Wiering. The meertens tune collections, December 2014.

[KM09] Alexis Kirke and Eduardo Reck Miranda. A survey of computer systems for expressive music performance. *ACM Computing Surveys (CSUR)*, 42(1):3, 2009.

[Kri07] David Kriesel. *A Brief Introduction to Neural Networks*. 2007.

[Lan89] Peter Langston. Six techniques for algorithmic music composition. In *15th International Computer Music Conference (ICMC), Columbus, Ohio, November*, pages 2–5. Citeseer, 1989.

[LK00] Mikael Laurson and Mika Kuuskankare. Towards idiomatic instrumental writing: A constraint based approach. In *Proceedings of the 2nd Annual Symposium on Systems Research in the Arts*, 2000.

[MM95] Eduardo Morales and Roberto Morales. Learning musical rules. In *Proceedings of the International Joint Conference on Artificial Inteligence*, 1995.

[MMBL12] Robert M MacCallum, Matthias Mauch, Austin Burt, and Armand M Leroi. Evolution of music by public choice. *Proceedings of the National Academy of Sciences*, 109(30):12081–12086, 2012.

[Moo72] James Anderson Moorer. Music and computer composition. *Communications of the ACM*, 15(2):104–113, February 1972.

[Moz98] W. A. Mozart. *Instruction - To compose without the least knowledge of Music so much Countrydances as one pleases, by throwing a certain Number with two Dice.* Simrock, ca. 1798.

[Mus14] Deutsches Musikinformationszentrum. Laienmusizieren in zahlen - ergebnisse bundesweiter studien und bevölkerungsumfragen. 10 2014.

[MW11] Daniel Müllensiefen and Geraint Wiggins. *Systematic Musicology Empirical and Theoretical Studies*, chapter Polynomial Functions as a Representation of Melodic Phrase Contour, pages 63–88. Peter Lang, 1 edition, 2011.

[PGMC97] Francisco C Pereira, Carlos Fernando Almeida Grilo, Luís Macedo, and Fernando Amílcar Bandeira Cardoso. Composing music with case-based reasoning. In *International Conference on Computational Models of Creative Cognition*, 1997.

[Rad74] Gary M Rader. A method for composing simple traditional music by computer. *Communications of the ACM*, 17(11):631–638, 1974.

[RP98] Rafael Ramirez and Julio Peralta. A constraint-based melody harmonizer. In *Proceedings of the Workshop on Constraints for Artistic Applications*, 1998.

[Sch93] Stephan M Schwanauer. *Machine models of music*, chapter A learning machine for tonal composition, pages 511–532. MIT Press, 1993.

[Sch99] Mark A. Schmuckler. Testing models of melodic contour similarity. *Music Perception: An Interdisciplinary Journal*, 16(3):295–326, 1999.

[Sch15] Jürgen Schmidhuber. Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117, 2015.

[Spa99] Randall Richard Spangler. *Rule-based analysis and generation of music.* PhD thesis, California Institute of Technology, 1999.

[Sta11] Statistische Ämter des Bundes und der Länder. Zensus 2011, 2011.

[WJ63] Joe H Ward Jr. Hierarchical grouping to optimize an objective function. *Journal of the American statistical association*, 58(301):236–244, 1963.

[WK07] Adam Wead and Ian Knopke. A computer-based implementation of basso continuo rules for figured bass realizations. In *ICMC*, pages 188–191, 2007.

[Wri05] Kevin B. Wright. Researching internet-based populations: Advantages and disadvantages of online survey research, online questionnaire authoring software packages, and web survey services. *Journal of Computer-Mediated Communication*, 10(3), 2005.

[Zim01] Detlev Zimmermann. Modelling musical structures. *Constraints*, 6(1):53–83, 2001.

Appendices

A. Screen shots of the conducted listening test

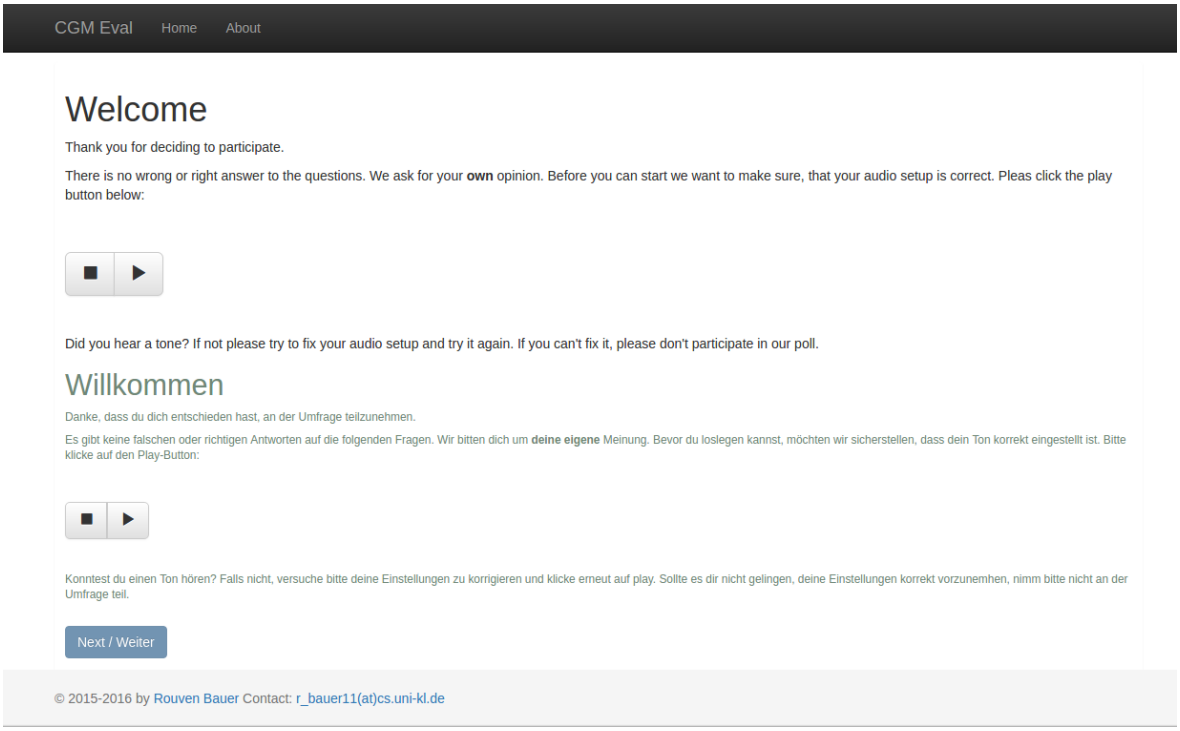
CGM Eval
Home
About
Welcome
Thank you for deciding to participate.
There is no wrong or right answer to the questions. We ask for your own opinion. Before you can start we want to make sure, that your audio setup is correct. Pleas click the play button below:
Did you hear a tone? If not please try to fix your audio setup and try it again. If you can't fix it, please don't participate in our poll.
Willkommen
Danke, dass du dich entschieden hast, an der Umfrage teilzunehmen.
Es gibt keine falschen oder richtigen Antworten auf die folgenden Fragen. Wir bitten dich um deine eigene Meinung. Bevor du loslegen kannst, möchten wir sicherstellen, dass dein Ton korrekt eingestellt ist. Bitte klicke auf den Play-Button:
Konntest du einen Ton hören? Falls nicht, versuche bitte deine Einstellungen zu korrigieren und klicke erneut auf play. Sollte es dir nicht gelingen, deine Einstellungen korrekt vorzunemhen, nimm bitte nicht an der Umfrage teil.
Next / Weiter
© 2015-2016 by Rouven Bauer Contact: r_bauer11(at)cs.uni-kl.de

Screen shot of test stimulus of conducted listening test.

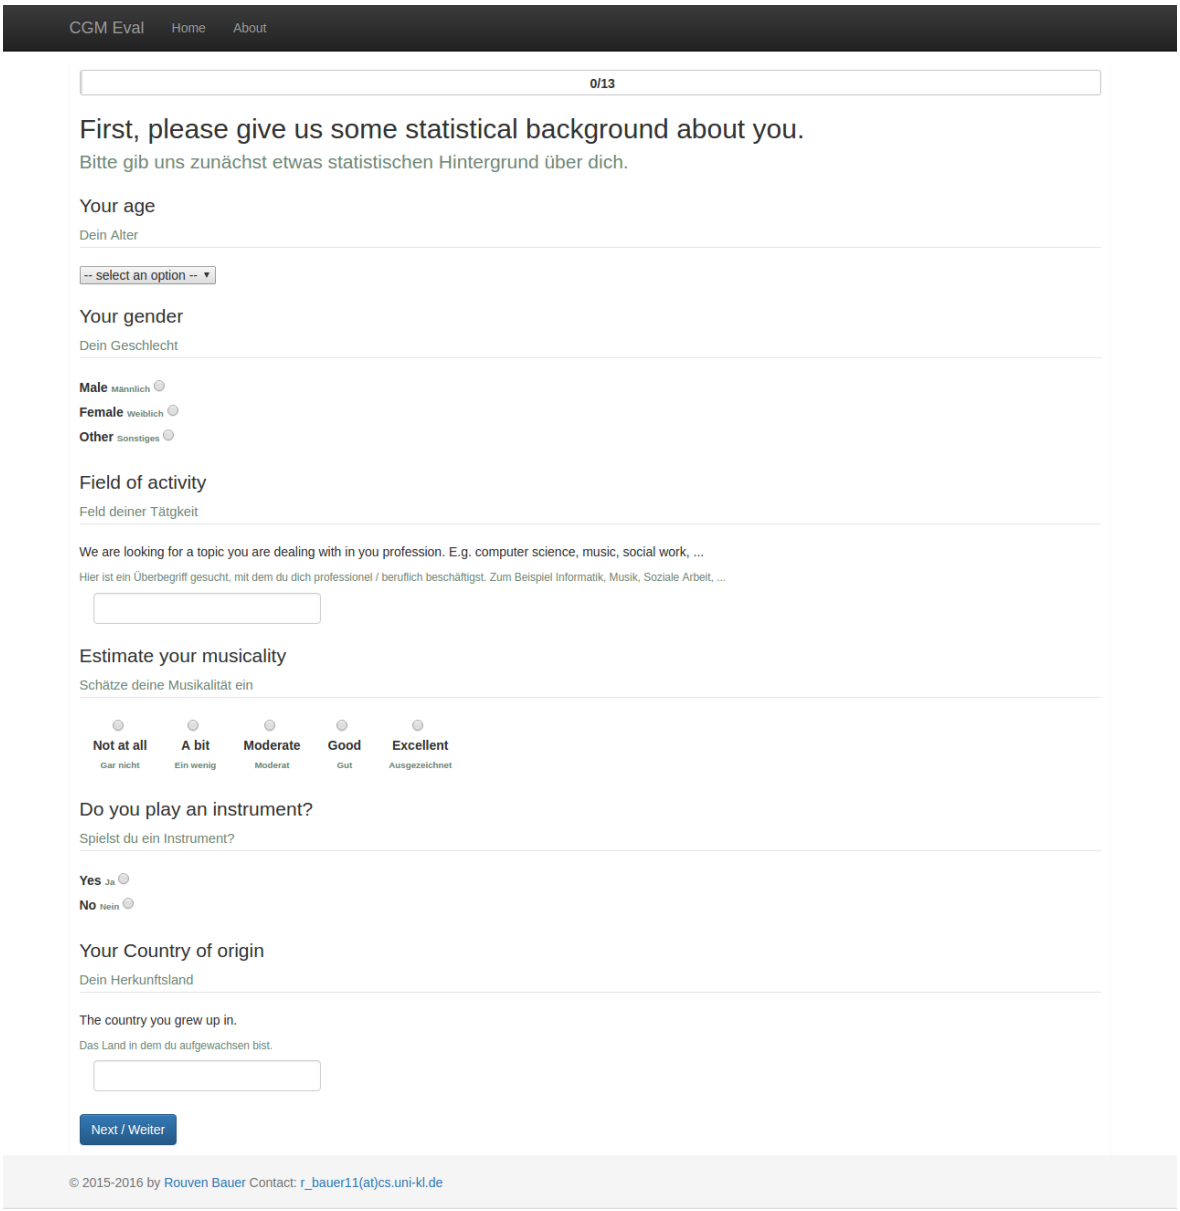

Screen shot of personal questions of conducted listening test.

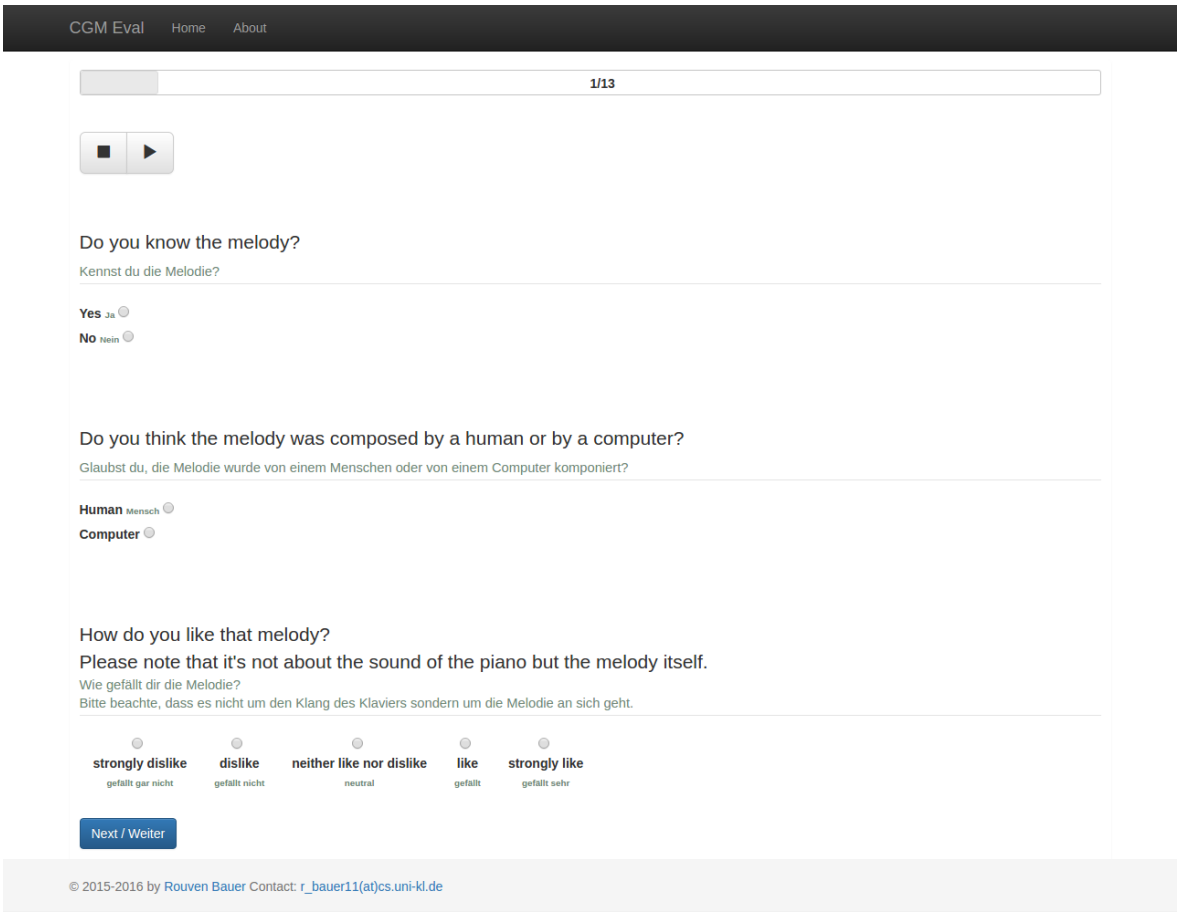
CGM Eval
Home
About
1/13
Do you know the melody?
Kennst du die Melodie?
Yes Ja
No Nein
Do you think the melody was composed by a human or by a computer?
Glaubst du, die Melodie wurde von einem Menschen oder von einem Computer komponiert?
Human Mensch
Computer
How do you like that melody?
Please note that it's not about the sound of the piano but the melody itself.
Wie gefällt dir die Melodie?
Bitte beachte, dass es nicht um den Klang des Klaviers sondern um die Melodie an sich geht.
strongly dislike
gefällt gar nicht
dislike
gefällt nicht
neither like nor dislike
neutral
like
gefällt
strongly like
gefällt sehr
Next / Weiter
© 2015-2016 by Rouven Bauer Contact: r_bauer11(at)cs.uni-kl.de

Screen shot of melody questions of conducted listening test. Each participant answered twelve of them.

B. Composition results with one training melody

Note that the training melody was transposed to c major before learning from it and composing new melodies. Keep that in mind when comparing the input with the outputs. KATH on the other hand, preserves the key of the training melodies.

B.1. Melody 1
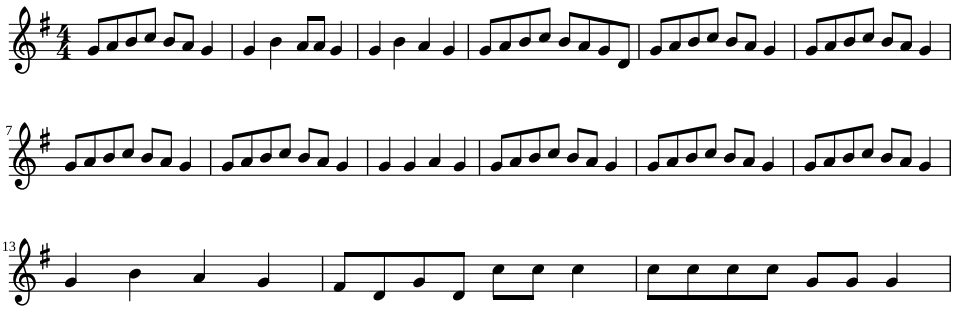
(a) Training melody
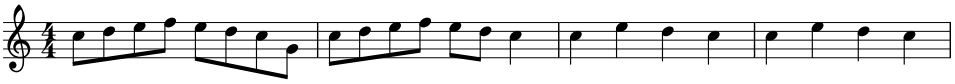
(b) composed by our algorithm
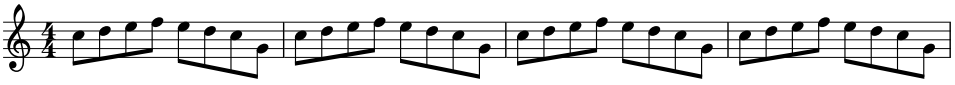
(c) composed by our algorithm

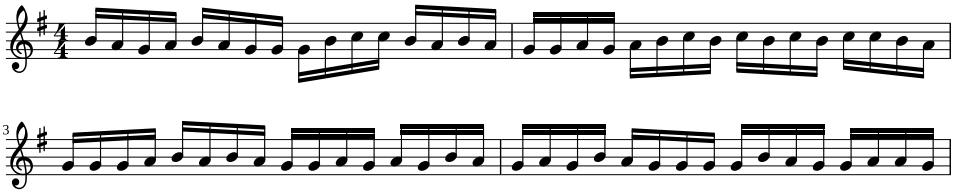

(d) composed by KATH

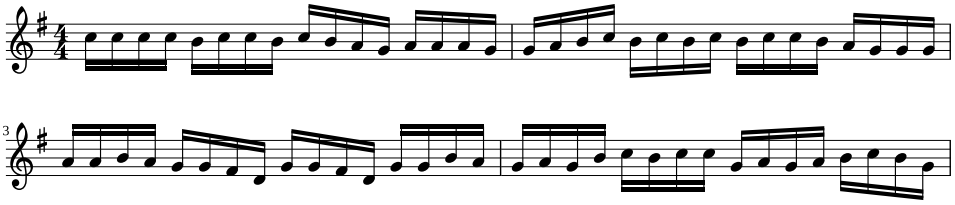

(e) composed by KATH

B.2. Melody 2
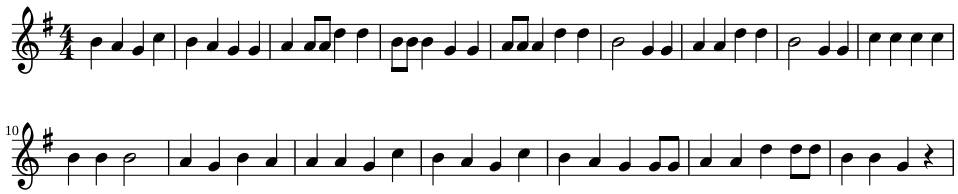
(f) Training melody
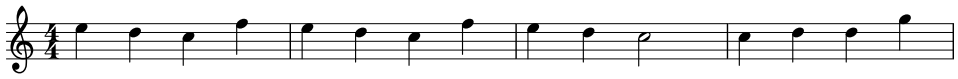
(g) Algorithmically composed melody a
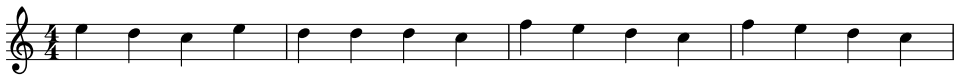
(h) composed by our algorithm

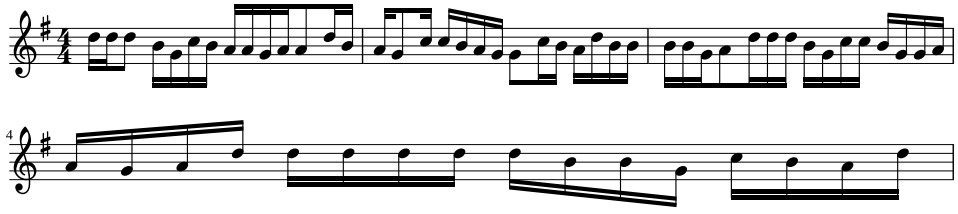

(i) composed by KATH

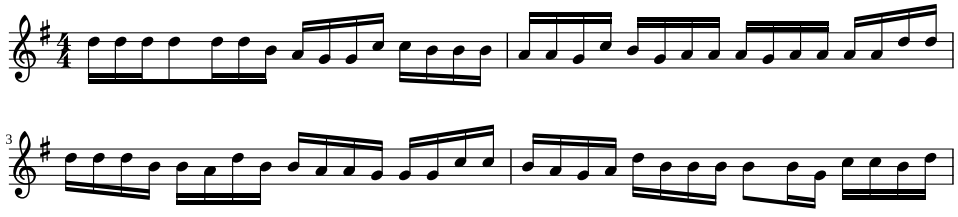

(j) composed by KATH

B.3. Melody 3

You might notice that both melodies generated by our algorithm are the same. Since a lot of randomness is involved in the composition process this is (depending on the width of the search tree) unlikely but possible. We did not filter those occurrences out because we did not want to modify the results in any matter. We actually didn't have a look at the output before starting the listening test.
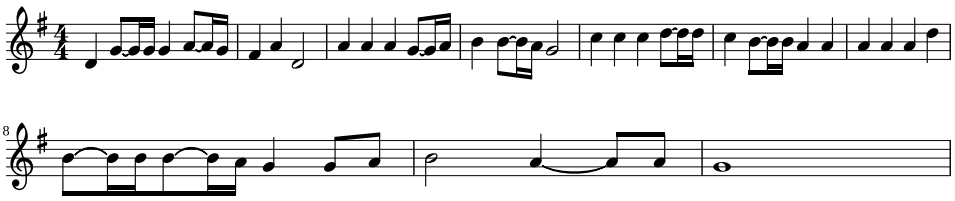
(k) Training melody
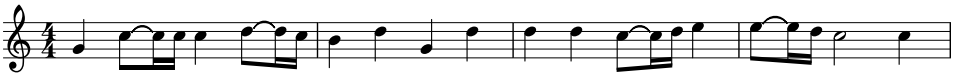
(l) composed by our algorithm
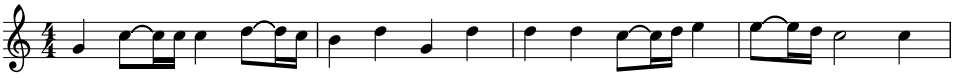
(m) composed by our algorithm

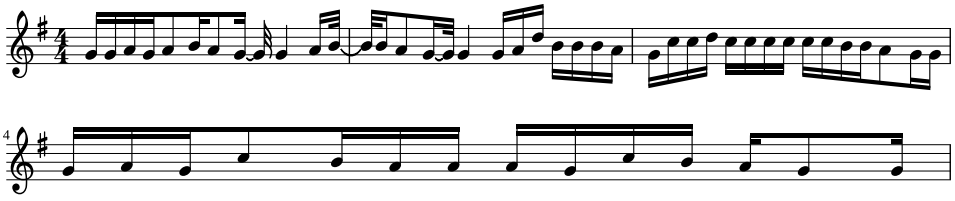

(n) composed by KATH
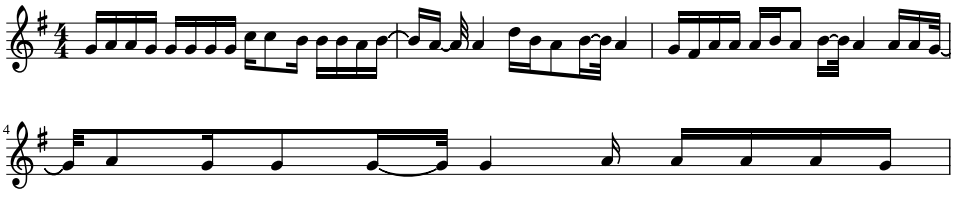
(o) composed by KATH

B.4. Melody 4
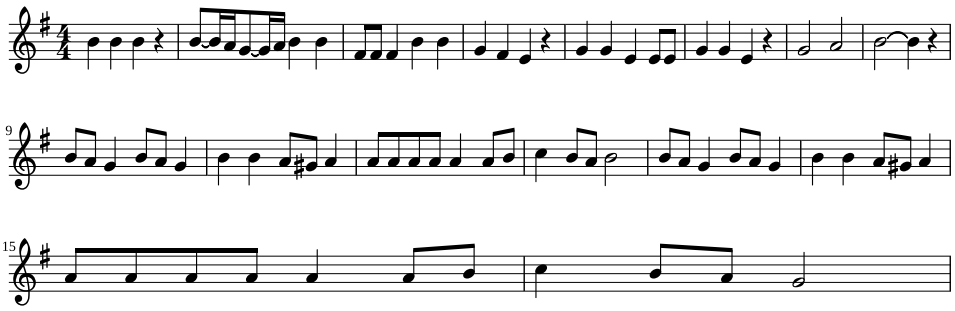
(p) Training melody
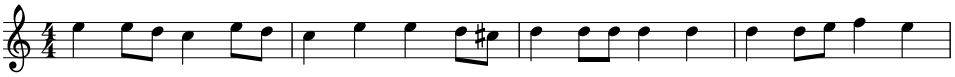
(q) composed by our algorithm
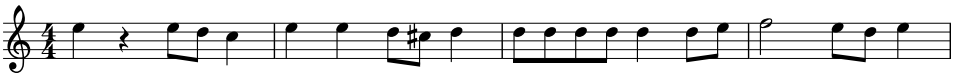
(r) composed by our algorithm

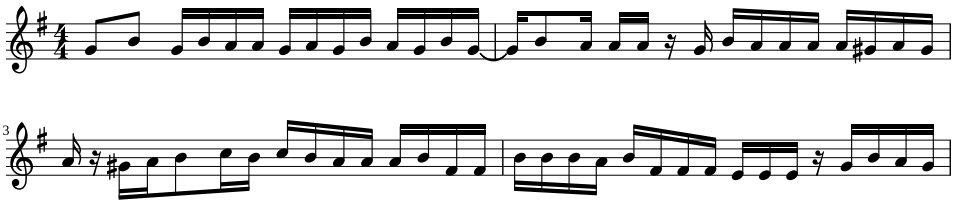

(s) composed by KATH
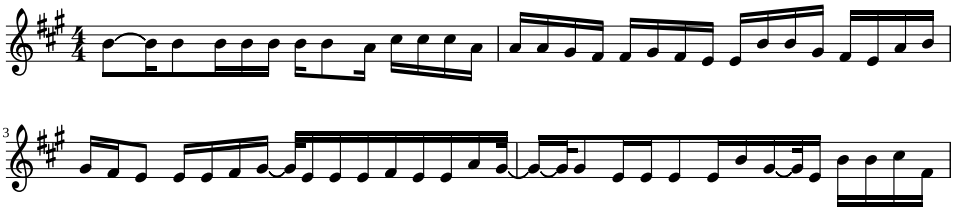
(t) composed by KATH

B.5. Melody 5
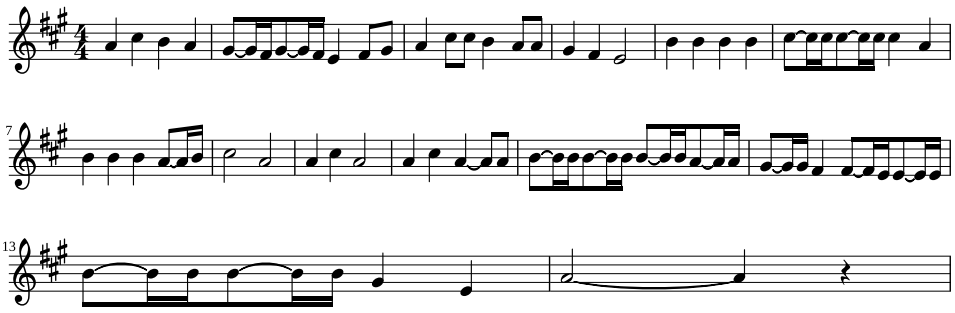
(u) Training melody
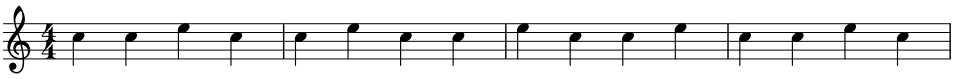
(v) composed by our algorithm
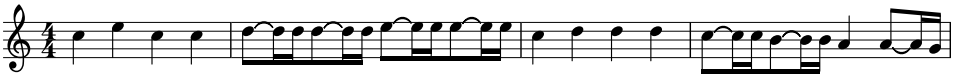
(w) composed by our algorithm
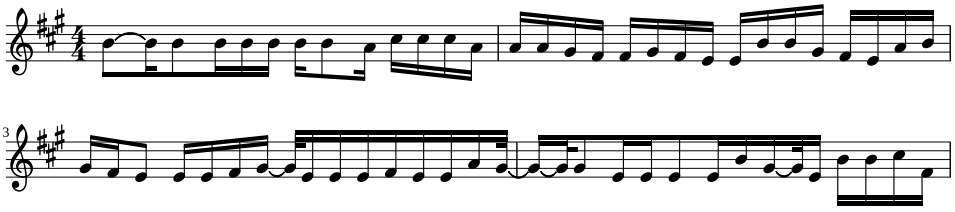
(x) composed by KATH
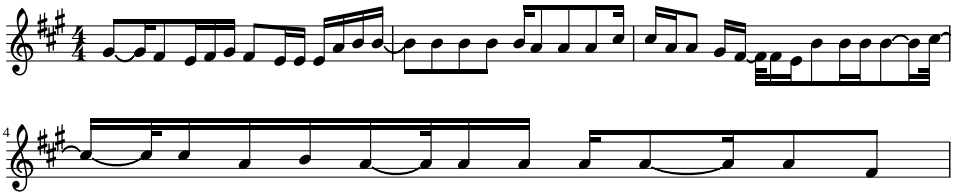
(y) composed by KATH